arXiv:1404.6835v1 [cs.DS] 27 Apr 2014

Bypassing Erdős' Girth Conjecture: Hybrid Stretch and Sourcewise Spanners

Merav Parter

The Weizmann Institute of Science, Rehovot, Israel. merav.parter@weizmann.ac.il *

Abstract. An (α, β)-spanner of an n-vertex graph $G = (V, E)$ is a subgraph H of G satisfying that $\text{dist}(u, v, H) \leq \alpha \cdot \text{dist}(u, v, G) + \beta$ for every pair $(u, v) \in V \times V$, where $\text{dist}(u, v, G')$ denotes the distance between u and v in $G' \subseteq G$. It is known that for every integer $k \geq 1$, every graph G has a polynomially constructible $(2k-1, 0)$-spanner of size $O(n^{1+1/k})$. This size-stretch bound is essentially optimal by the girth conjecture. Yet, it is important to note that any argument based on the girth only applies to *adjacent vertices*. It is therefore intriguing to ask if one can "bypass" the conjecture by settling for a multiplicative stretch of $2k-1$ only for *neighboring* vertex pairs, while maintaining a strictly *better* multiplicative stretch for the rest of the pairs. We answer this question in the affirmative and introduce the notion of *k-hybrid spanners*, in which non neighboring vertex pairs enjoy a *multiplicative* k-stretch and the neighboring vertex pairs enjoy a *multiplicative* $(2k-1)$ stretch (hence, tight by the conjecture). We show that for every unweighted n-vertex graph G with m edges, there is a (polynomially constructible) k-hybrid spanner with $O(k^2 \cdot n^{1+1/k})$ edges. This should be compared against the current best (α, β) spanner construction of [5] that obtains $(k, k-1)$ stretch with $O(k \cdot n^{1+1/k})$ edges. An alternative natural approach to bypass the girth conjecture is to allow ourself to take care only of a subset of pairs $S \times V$ for a given subset of vertices $S \subseteq V$ referred to here as *sources*. Spanners in which the distances in $S \times V$ are bounded are referred to as *sourcewise spanners*. Several constructions for this variant are provided (e.g., multiplicative sourcewise spanners, additive sourcewise spanners and more).

1 Introduction

1.1 Motivation

Graph spanners are sparse subgraphs that faithfully preserve the pairwise distances of a given graph and provide the underlying graph structure in communication networks, robotics, distributed systems and more [26]. The notion

* Recipient of the Google European Fellowship in distributed computing; research supported in part by this Fellowship. Supported in part by the Israel Science Foundation (grant 894/09), United States-Israel Binational Science Foundation (grant 2008348), Israel Ministry of Science and Technology (infrastructures grant), and Citi Foundation.

of graph spanners was introduced in [24,25] and have been studied extensively since. Spanners have a wide range of applications from distance oracles [30,8], labeling schemes [9] and routing [13] to solving linear systems [17] and spectral sparsification [19].

Given an undirected unweighted n-vertex graph $G = (V, E)$, a subgraph H of G is said to be a k-spanner if for every pair of vertices $(u, v) \in V \times V$ it holds that $\text{dist}(u, v, H) \leq k \cdot \text{dist}(u, v, G)$. It is well known that one can efficiently construct a $(2k-1)$-spanner with $O(n^{1+1/k})$ edges, even for weighted graphs [4,6]. This size-stretch ratio is conjectured to be tight based on the girth[1] conjecture of Erdős [18], which says that there exist graphs with $\Omega(n^{1+1/k})$ edges and girth $2k+1$. If one removes an edge in such a graph, the distance between the edge endpoints increases from 1 to $2k$, implying that any α-spanner for $\alpha \leq 2k-1$ has $\Omega(n^{1+1/k})$ edges. This conjecture has been resolved for the special cases of $k = 1, 2, 3, 5$ [32].

Although the girth conjecture exactly characterizes the optimal tradeoff between sparseness and multiplicative stretch, it applies only to adjacent vertices (i.e., removing an edge (u, v) from a large cycle causes distortion to the edge endpoints). Indeed, Elkin and Peleg [15] showed that the girth bound (on multiplicative distortion) fails to hold even for vertices at distance 2. This limitation of the girth argument motivated distinguishing between *nearby* vertex pairs and "sufficiently distant" vertex pairs. This gave raise to the development of (α, β)-spanners which distort distances in G up to a multiplicative factor of α and an additive term β [15]. Formally, for an unweighted undirected graph $G = (V, E)$, a subgraph H of G is an (α, β)-spanner iff $\text{dist}(u, v, H) \leq \alpha \cdot \text{dist}(u, v, G) + \beta$ for every $u, v \in V$. Note, that an (α, β)-spanner makes an implicit distinction between nearby vertex pairs and sufficiently distant vertex pairs. In particular, for "sufficiently distant" vertex pairs the (α, β)-spanner behaves similar to a pure multiplicative spanner, whereas for the remaining vertex pairs, the spanner behaves similar to an additive spanner [21]. The setting of (α, β)-spanners has been widely studied for various distortion-sparseness tradeoffs [16,31,15,5]. For example, [15] gave a construction for $(k-1, 2k-O(1))$-spanners with size $O(k \cdot n^{1+1/k})$, with a number of refinements for short distances, and showed that for any $k \geq 2$ and $\epsilon > 0$, there exist $(1+\epsilon, \beta)$-spanners with size $O(\beta \cdot n^{1+1/k})$, where β depends on ϵ and k but independent on n, implying that the size can be driven close to linear in n and the multiplicative stretch close to 1, at the cost of a large additive term in the stretch. Thorup and Zwick designed $(1+\epsilon, \beta)$-spanners with $O(k \cdot n^{1+1/k})$ edges, with a multiplicative distortion that tends to 1 as the distance increases [31].

The best (α, β) spanner construction is due to [5] which achieves stretch of $(k, k-1)$ with $O(k \cdot n^{1+1/k})$ edges, hence providing multiplicative stretch $2k-1$ for neighboring vertices (which is the best possible by Erdős' conjecture) and a multiplicative stretch at most $3k/2$ for the remaining pairs.

Although (α, β)-spanners make an (implicit) distinction between "close" and "distant" vertex pairs, as the girth argument holds only for vertices at dis-

[1] The girth is the smallest cycle length.

tance 1, it seems that a tighter bound on the behavior of spanners may be obtained. In particular, it seems plausible that the multiplicative factor of k using $O(n^{1+1/k})$ edges, is not entirely unavoidable for non-neighboring vertex pairs, while providing multiplicative stretch of $2k-1$ for the neighboring vertex pairs. The current paper confirms this intuition by introducing the notion of k-hybrid spanners, namely, subgraphs $H \subseteq G$ that obtain multiplicative stretch $2k-1$ for neighboring vertices, i.e., $\text{dist}(u,v,H) \leq (2k-1)\cdot\text{dist}(u,v,G)$ for every $(u,v) \in E(G)$ and multiplicative stretch k for the remaining vertex pairs, i.e., $\text{dist}(u,v,H) \leq k \cdot \text{dist}(u,v,G)$ for every $(u,v) \notin E(G)$. Hence, hybrid spanners seem to pinpoint the minimum possible relaxation of the stretch requirement in spanners graphs so that the girth conjecture lower bound can be bypassed. The presented k-hybrid spanner with $O(k^2 \cdot n^{1+1/k})$ edges can be contrasted with several existing spanner constructions, e.g, k-spanners with $O(n^{1+2/(k+1)})$ edges (in which multiplicative stretch k is guaranteed also to neighboring pairs), the $\Omega(k \cdot n^{1+1/k})$ lower-bound graph construction for $(2k-1)$-additive spanners, and to the $(k, k-1)$ spanner construction of [5] with $O(k^2 \cdot n^{1+1/k})$ edges.

An alternative approach to bypass the conjecture is by focusing on a subset of pairs in $V \times V$. Following [10,27,12,20], we relax the requirement that small stretch in the subgraph must be guaranteed for *every* vertex pair from $V \times V$. Instead, we require it to hold only for pairs of vertices from a subset of $V \times V$. Specifically, given a subset of vertices $S \subseteq V$, referred to here as *sources*, our spanner H aims to bound only the distances between pairs of vertices from $S \times V$. For any other pair outside $S \times V$, the stretch in H can be arbitrary.

On the lower bound side, Woodruff [33] proved, independently of the Erdős' conjecture, the existence of graphs for which any spanner of size $\Omega(k^{-1}n^{1+1/k})$ has an additive stretch of at least $2k-1$. Although sourcewise additive spanners have been studied by [27,12,20], currently there are no known lower bound constructions for this variant. We generalize Woodruff's construction to the sourcewise setting, providing a graph construction whose size has a smooth dependence with the number of sources.

1.2 Related Works

The notion of a sparse subgraph that preserves distances only for a subset of the $V \times V$ pairs has been initiated by Bollobás, Coopersmith and Elkin [9], who studied *pairwise preservers*, where the input is a graph $G = (V, E)$ along with a subset of vertex pairs $\mathcal{P} \subseteq V \times V$ and the problem is to construct a sparse subgraph H such that the $u-v$ distance for each $(u,v) \in \mathcal{P}$ is exactly preserved, i.e., $\text{dist}(u,v,H) = \text{dist}(u,v,G)$ for every $(u,v) \in \mathcal{P}$. They showed that one can construct a pairwise preserver with $O(\min\{|\mathcal{P}| \cdot \sqrt{n}, n \cdot \sqrt{|\mathcal{P}|}\})$ edges. At the end of their paper, they raised the question of constructing sparser subgraphs where distances between pairs in $\mathcal{P}$ are *approximately* preserved, or in other words, the problem of constructing sparse $\mathcal{P}$-spanners. Pettie [27] studied a certain type of $\mathcal{P}$-spanners, namely, additive sourcewise spanners. In this setting, one is given an unweighted graph $G = (V, E)$ and a subset of vertices $S \subseteq V$, termed as *sources*, whose size is conveniently parameterized to be $|S| = n^{\varepsilon}$, and the goal is

to construct a sparse spanner H that maintains an additive approximation for the $S \times V$ distances. He showed a construction of $O(\log n)$-additive sourcewise spanners of size $O(n^{1+\varepsilon/2})$. Cygan et al. recently showed a stretch-size bound for $2k$-additive sourcewise spanners with $O(n^{1+(\varepsilon k+1)/(2k+1)})$ edges. The specific case of $k = 1$ has been studied recently by [20], providing a 2-additive sourcewise spanner with $\widetilde{O}(n^{5/4+\varepsilon/4})$ edges where $|S| = n^{\varepsilon}$.

Upper bounds for spanners with constant stretch are currently known for but a few stretch values. A $(1, 2)$ spanner with $O(n^{3/2})$ edges is presented in [2], a $(1, 6)$ spanner with $O(n^{4/3})$ edges is presented in [5], and a $(1, 4)$ spanner with $O(n^{7/5})$ edges is presented in [11]. The latter two constructions use the *path-buying* strategy, which is adopted in our additive sourcewise construction. Dor et al. [14] considered additive emulators, which may contain additional (possibly weighted) edges. They showed a construction of 4-additive emulator with $O(n^{4/3})$ edges. Finally, a well known application of α-spanners is *approximate distance oracles* [30,23,8,7,22]. The sourcewise variant, namely, *sourcewise approximate distance oracle* was devised by [28]. For a given input graph $G = (V, E)$ and a source set $S \subseteq V$, [28] provides a construction of a distance oracle of size $O(n^{1+\varepsilon/k})$ where $\varepsilon = \log |S|/\log n$ such that given a distance query $(s, v) \in S \times V$ returns in $O(k)$ time a $(2k - 1)$ approximation to $\mathrm{dist}(s, v, G)$.

1.3 Contributions

In this paper we initiate the study of k-hybrid spanners which seems to pinpoint the minimal condition for bypassing Erdős' Girth Conjecture. In addition, we also study the sourcewise variant of multiplicative spanners, additive spanners and additive emulators. The main results are summarized below.

Theorem 1 (Hybrid spanners). *For every integer $k \geq 2$ and unweighted undirected n-vertex graph $G = (V, E)$, there exists a (polynomially constructible) subgraph of size $O(k^2 \cdot n^{1+1/k})$ that provides multiplicative stretch $2k - 1$ for every pair of neighboring vertices u and v and a multiplicative stretch k for the rest of the pairs. (By Erdős' conjecture, providing a multiplicative stretch of k for all the pairs requires $\Omega(n^{1+2/(k+1)})$ edges.)*

Theorem 2 (Sourcewise spanners). *For every integer $k \geq 2$, and an unweighted undirected n-vertex graph $G = (V, E)$ and for every subset of sources $S \subseteq V$ of size $|S| = O(n^{\varepsilon})$, there exists a (polynomially constructible) subgraph of size $O(k^2 \cdot n^{1+\varepsilon/k})$ that provides multiplicative stretch $2k - 1$ for every pair of neighboring vertices $(u, v) \in S \times V$ and a multiplicative stretch of $2k - 2$ for the rest of the pairs in $S \times V$. This subgraph is referred to here as* sourcewise spanner*.*

Theorem 3 (Lower bound for additive sourcewise spanners and emulators). *For every integer $k \in [2, O(\log n/\log\log n)]$ and $\varepsilon \in [0, 1]$, there exists an n-vertex graph $G = (V, E)$ and a subset of sources $S \subseteq V$ of size $|S| = O(n^{\varepsilon})$ such that any $(2k - 1)$-additive sourcewise spanner (i.e., subgraph that maintains a $(2k - 1)$-additive approximation for the $S \times V$ distances) has at least*

$\Omega(n^{1+\varepsilon/k}/k)$ *edges. The lower bound holds for additive emulators up to order* $O(k)$*. For* 2*-additive sourcewise emulators there is a matching upper bound.*

Theorem 4 (Upper bound for additive sourcewise spanners). *Let* $k \geq 1$ *be an integer. (1) For every unweighted undirected* n*-vertex graph* $G = (V, E)$ *and for every subset of sources* $S \subseteq V$*,* $|S| = O(n^{\varepsilon})$*, there exists a (polynomially constructible)* $2k$*-additive sourcewise spanner with* $\widetilde{O}(k \cdot n^{1+(\varepsilon \cdot k+1)/(2k+2)})$ *edges. (2) For* $|S| = \Omega(n^{2/3})$*, there exists a* 4*-additive sourcewise spanner with* $O(n^{1+\varepsilon/2})$ *edges (by the lower bound of Thm. 3, any* 3*-additive sourcewise spanner requires* $\Omega(n^{1+\varepsilon/2})$ *edges).*

The time complexities of all our upper bound construction are obviously polynomial; precise analysis is omitted from this extended abstract.

1.4 Preliminaries

We consider the following graph structures.

(α, β)-spanners. For a graph $G = (V, E)$, the subgraph $H \subseteq G$ is an (α, β)-spanner for G if for every $(u, v) \in V \times V$,

$$\mathrm{dist}(u, v, H) \leq \alpha \cdot \mathrm{dist}(u, v, G) + \beta \ . \tag{1}$$

$(\alpha, 0)$-spanners (resp., $(1, \beta)$-spanners) are referred to here as α-spanners (resp., β-additive spanners).

Hybrid spanners. Given a graph $G = (V, E)$, a subgraph $H \subseteq G$ is a k-hybrid spanner iff for every $(u, v) \in V \times V$ it holds that

$$\mathrm{dist}(u, v, H) \leq \begin{cases} (2k-1) \cdot \mathrm{dist}(u, v, G), & \text{if } (u, v) \in E(G); \\ k \cdot \mathrm{dist}(u, v, G), & \text{otherwise.} \end{cases} \tag{2}$$

Sourcewise spanners. Given an unweighted graph $G = (V, E)$ and a subset of vertices $S \subseteq V$, a subgraph $H \subseteq G$ is an (α, β, S)*-spanner* iff Eq. (1) is satisfied for every $\langle s, v \rangle \in S \times V$. When $\beta = 0$ (resp., $\alpha = 1$), H is denoted by (α, S)*-sourcewise spanner* (resp., (β, S)*-additive sourcewise spanner*).

Emulators. Given an unweighted graph $G = (V, E)$, a weighted graph $H = (V, F)$ is an (α, β)*-emulator* of G iff $\mathrm{dist}(u, v, G) \leq \mathrm{dist}(u, v, H) \leq \alpha \cdot \mathrm{dist}(u, v, G) + \beta$ for every $(u, v) \in V \times V$. $(1, \beta)$-emulators are referred to here as β-additive emulators. For a given subset of sources $S \subseteq V$, the graph $H = (V, F)$ is a (β, S)*-additive sourcewise emulator* if the $S \times V$ distances are bounded in H by an additive stretch of β.

1.5 Notation

For a subgraph $G' = (V', E') \subseteq G$ (where $V' \subseteq V$ and $E' \subseteq E$) and a pair of vertices $u, v \in V'$, let $\mathrm{dist}(u, v, G')$ denote the shortest-path distance in edges between u and v in G'. Let $\Gamma(v, G) = \{u \mid (u, v) \in E(G)\}$ be the set of

neighbors of v in G. For a subgraph $G' \subseteq G$, let $|G'| = |E(G')|$ denote the number of edges in G'. For a path $P = [v_1, \ldots, v_k]$, let $P[v_i, v_j]$ be the subpath of P from v_i to v_j. For paths P_1 and P_2, let $P_1 \circ P_2$ denote the path obtained by concatenating P_2 to P_1. Let $SP(s, v_i, G')$ be the set of $s - v_i$ shortest-paths in G'. When G' is the input graph G, let $\pi(x, y) \in SP(x, y, G)$ denote some arbitrary $x - y$ shortest path in G, hence $|\pi(x, y)| = \text{dist}(x, y, G)$. For a subset $V' \subseteq V$, let $\text{dist}(u, V', G) = \min_{u' \in V'} \text{dist}(u, u', G)$. Similarly, for subsets $V_1, V_2 \subseteq V$, $\text{dist}(V_1, V_2, G) = \min_{v_1 \in V_1, v_2 \in V_2} \text{dist}(v_1, v_2, G)$. When the graph G is clear from the context, we may omit it and simply write $\Gamma(u), \text{dist}(u, v), \text{dist}(u, V')$ and $\text{dist}(V_1, V_2)$.

A clustering $\mathcal{C} = \{C_1, \ldots, C_\ell\}$ is a collection of disjoint subsets of vertices, i.e., $C_i \subseteq V$ for every $C_i \in \mathcal{C}$ and $C_i \cap C_j = \emptyset$ for every $C_i, C_j \in \mathcal{C}$. Note that a clustering is not necessarily a partition of V, i.e., it is not required that $\bigcup_i C_i = V$. A cluster $C \in \mathcal{C}$ is said to be *connected* in G if the induced graph $G[C]$ is connected. For clusters C and C', let $E(C, C') = (C \times C') \cap E(G)$ be the set of edges between C and C' in G. For notational simplicity, let $E(v, C) = E(\{v\}, C)$. A vertex v is *incident* to a cluster C if $E(v, C) \neq \emptyset$. In a similar manner, two clusters C and C' are adjacent to each other if $E(C, C') \neq \emptyset$.

Organization. We start with upper bounds. Sec. 2 describes the construction of k-hybrid spanners. Sec. 3.1, presents the construction of (α, S) sourcewise spanners. Then, Sec. 3.2 presents a lower bound construction for (β, S) sourcewise additive spanners and emulators. Finally, Sec. 3.3 provides an upper bound for $(2k, S)$-additive sourcewise spanners for general values of k. In addition, it provides a tight construction for $(2, S)$-additive sourcewise emulators.

2 Hybrid spanners

In this section, we establish Thm. 1. For clarity of presentation, we describe a randomized construction whose output spanner has $O(k^2 \cdot n^{1+1/k})$ edges in expectation. Using the techniques of [5], this construction can be derandomized with the same bound on the number of edges.

Theorem 5. *Let $k \geq 2$ be an integer. For every unweighted n-vertex graph $G = (V, E)$ with m edges, a k-hybrid spanner $H \subseteq G$ with $O(k^2 \cdot n^{1+1/k})$ edges can be constructed in $O(k^2 \cdot m)$ time.*

The algorithm. We begin by describing a basic procedure Cluster, slightly adapted from [5], that serves as a building block in our constructions. For an input unweighted graph $G = (V, E)$, a stretch parameter k and a density parameter μ, Algorithm Cluster iteratively constructs a sequence of $k+1$ clusterings $\mathcal{C}_0, \ldots, \mathcal{C}_k$ and a clustering graph $H_k \subseteq G$. Each clustering $\mathcal{C}_\tau$ consists of $m_\tau = n^{1-\tau \cdot \mu}$ disjoint subsets of vertices, $C_\tau = \{C_1^\tau, \ldots, C_{m_\tau}^\tau\}$. Each cluster $C_j^\tau \in \mathcal{C}_\tau$ is connected and has a *cluster center* z_j satisfying that $\text{dist}(u, z_j, G) \leq \tau$ for every $u \in C_j^\tau$. Denote the set of cluster centers of $\mathcal{C}_\tau$ by Z_τ. These cluster centers correspond to a sequence of samples taken from V with decreasing densities

where $V = Z_0 \supseteq Z_1 \supseteq \ldots \supseteq Z_k$. On a high level, at each iteration τ, a clustering of radius-τ clusters is constructed and its shortest-path spanning forest (spanning all the vertices in the clusters), as well as an additional subset of edges Q_τ adjacent to unclustered vertices, are chosen to be added to the spanner H_τ. We now describe the algorithm $\mathsf{Cluster}(G, k, \mu)$ in detail. Assume some ordering on the vertices $V = \{v_1, \ldots, v_n\}$. Initially, the cluster centers are $Z_0 = V = \{v_1, \ldots, v_n\}$, where each vertex forms its own cluster of radius 0, hence $\mathcal{C}_0 = \{\{v\} \mid v \in V\}$ and the spanner is initiated to $H_0 = \emptyset$. At iteration $\tau \geq 1$, a clustering $\mathcal{C}_\tau$ is defined based on the cluster centers $Z_{\tau-1}$ of the previous iteration. Let $Z_\tau \subseteq Z_{\tau-1}$ be a sample of $m_\tau = O(n^{1-\tau\cdot\mu})$ vertices chosen uniformly at random from $Z_{\tau-1}$. The clustering $\mathcal{C}_\tau$ is obtained by assigning every vertex u that satisfies $\mathrm{dist}(u, Z_\tau, G) \leq \tau$ to its closest cluster center $z \in Z_\tau$, i.e., such that $\mathrm{dist}(u, z, G) = \mathrm{dist}(u, Z_\tau, G)$. If there are several cluster centers in Z_τ at distance $\mathrm{dist}(u, Z_\tau, G)$ from u, then the closest center with the minimal index is chosen.

Formally, for a vertex v and subset of vertices B, let $\mathtt{nearest}(u, B)$ be the closest vertex to u in B where ties are determined by the indices, i.e., letting $B' = \{u_1, \ldots, u_\ell\} \subseteq B$ be the set of closest vertices to u in B, namely, satisfying that $\mathrm{dist}(v, u_1) = ... = \mathrm{dist}(v, u_\ell) = \mathrm{dist}(v, B)$, then $\mathtt{nearest}(u, B) \in B'$ and has the minimal index in B'. Then u is assigned to the cluster of the center $\mathtt{nearest}(u, Z_\tau)$. Add to H_τ the forest F_τ consisting of the radius-τ spanning tree of each $C \in \mathcal{C}_\tau$. Note that the definition of the clusters immediately implies their connectivity. Next, an edge set Q_τ adjacent to unclustered vertices is added to H_τ as follows. Let Δ_τ denote the set of vertices that occur in each of the clusterings $\mathcal{C}_0, \ldots, \mathcal{C}_{\tau-1}$ but do not occur in $\mathcal{C}_\tau$. (Observe that such a vertex may re-appear again in some future clusterings.) Formally, let $\widehat{V}_\tau = \bigcup_{C \in \mathcal{C}_\tau} C$ be the set of vertices that occur in some cluster in the clustering $\mathcal{C}_\tau$. Then, $\Delta_\tau = \left(\bigcap_{j=0}^{\tau-1} \widehat{V}_j\right) \setminus \widehat{V}_\tau$. Note that by this definition, each vertex belongs to at most one set Δ_τ. Hence we have:

Observation 6 *The sets Δ_τ are disjoint.*

For every vertex $v \in \Delta_\tau$ and every cluster $C \in \mathcal{C}_{\tau-1}$ that is adjacent to v, pick one vertex $u \in C$ adjacent to v and add the edge (u, v) to Q_τ. (In other words, an edge (u, v) is *not* added to Q_τ for $v \in \Delta_\tau$ if either $u \notin \widehat{V}_{\tau-1}$ or an edge (u', v) was added to Q_τ where u' and u are in the same cluster $C \in \mathcal{C}_{\tau-1}$.) Then add Q_τ to H_τ. This completes the description of Algorithm $\mathsf{Cluster}$; a pseudocode is given below.

Algorithm $\mathsf{Cluster}(G, k, \mu)$.

(T1) Let $H_0 = \emptyset$ and $Z_0 = n$. Select a sample Z_τ uniformly at random from $Z_{\tau-1}$ with probability $n^{-\mu}$ for $\tau = 1$ to k (if $\mu = 1$ and $\tau = k$, set $Z_k = \emptyset$).

(T2) For $\tau = 1$ to k, define the clustering $\mathcal{C}_\tau$ by adding the τ-radius neighborhood for all cluster centers Z_τ, i.e., every $u \in V$ satisfying $\mathrm{dist}(u, Z_\tau) \leq \tau$ is connected to $\mathtt{nearest}(u, Z_\tau)$. Let F_τ denote the τ-radius neighborhood forest corresponding to $\mathcal{C}_\tau$.

(T3) For every vertex $v \in \Delta_\tau$ that was unclustered in the clustering $\mathcal{C}_\tau$ for the first time, let $e(v, C)$ be an arbitrary edge from $E(v, C)$ for every $C \in \mathcal{C}_{\tau-1}$.
(T4) $H_\tau = H_{\tau-1} \cup F_\tau \cup \{e(v, C) \mid v \in \Delta_\tau, C \in \mathcal{C}_{\tau-1}\}$.

The first step of Algorithm ConsHybrid applies Algorithm Cluster(G, k, μ) for $\mu = 1/k$, resulting in the graph H_k. Note that by Thm. 3.1 of [5], H_k is a $(2k-1)$ spanner. Hence, the stretch for neighboring vertices is $(2k-1)$ as required. We now add two edge sets to H_k in order to provide a multiplicative stretch k for the remaining pairs. Let

$$t = \lfloor k/2 \rfloor \quad \text{and} \quad t' = k - 1 - t, \tag{3}$$

Note that $t' = t$ when k is odd and $t' = t-1$ when k is even, so in general $t' \leq t$.

The algorithm considers the collection of $Z_{t'} \times Z_t$ shortest paths $\mathcal{P} = \{\pi(z_i, z_j) \mid z_i \in Z_{t'} \text{ and } z_j \in Z_t\}$. Starting with $H = H_k$, for each path $\pi(z_i, z_j) \in \mathcal{P}$, it adds to H the ℓ_t last edges of $\pi(z_i, z_j)$ (closest to z_i), where

$$\ell_t = 7t + 8t^2 \ . \tag{4}$$

For every pair of clusters C_1, C_2, let $\pi(C_1, C_2)$ denote the shortest path in G between some closest vertices $u_1 \in C_1$ and $u_2 \in C_2$ (i.e., $\text{dist}(C_1, C_2, G) = \text{dist}(u_1, u_2, G)$). For every τ from 0 to $k-1$, and for every pair $C_1 \in \mathcal{C}_\tau$ and $C_2 \in \mathcal{C}_{k-1-\tau}$, the algorithm adds to H, the ℓ last edges of $\pi(C_1, C_2)$, where $\ell = \ell_t$ for $\tau \in \{t', t\}$ and $\ell = 2k-1$ otherwise. This completes the description of Algorithm ConsHybrid, whose summary is given below.

Algorithm ConsHybrid.

(S1) Let $H_k =$ Cluster$(G, k, 1/k)$.
(S2) Let E_2 be the edge set containing the last ℓ_t edges of the path $\pi(z_i, z_j)$ for every $z_i \in Z_{t'}$ and $z_j \in Z_t$.
(S3) Let E_3 be the edges set containing, for every $\tau \in \{0, \ldots, k-1\}$, and for every $C_1 \in \mathcal{C}_\tau$ and $C_2 \in \mathcal{C}_{k-1-\tau}$, the last ℓ edges of the path $\pi(C_1, C_2)$ where $\ell = \ell_t$ for $\tau \in \{t', t\}$ and $\ell = 2k-1$ otherwise.
(S4) Let $H \leftarrow H_k \cup E_2 \cup E_3$.

Note that our algorithm bares some similarity to the $(k, k-1)$ construction of [5] but the analysis is different. The key difference between these two constructions is that in [5] only edges (i.e., shortest-path of length 1) are added between certain pairs of clusters. In contrast, in our construction, $O(k^2)$ edges are taken from each shortest-path connecting the close-most vertices coming from certain subset of clusters. This allows us to employ an inductive argument on the desired *purely* multiplicative stretch, without introducing an additional additive stretch term (for non-neighboring pairs) as in the $(k, k-1)$ construction of [5]. Specifically, by adding paths of length ℓ_t between center pairs in $Z_{t'} \times Z_t$, a much better stretch guarantee can be provided for (non-neighboring) $Z_{t'} \times Z_t$ pairs: a multiplicative stretch k plus a *negative* additive term. This additive term is then increased but in a controlled manner (due to step (S3)), resulting in a *zero* additive term for *any* non-neighboring vertex pair in $V \times V$.

Analysis We begin with size analysis. An edge $e \in E(G)$ is *missing* if it is not included in the spanner, i.e., $e \in E(G) \setminus E(H)$. We first bound the expected size of the partial spanner H_k and the final spanner H obtained at the end of phase (S1).

Lemma 1. *(1)* $\mathbb{E}(|H_k|) = O(n^{1+\mu} + k \cdot n)$.
(2) $\mathbb{E}(|H|) = O(k^2 \cdot n^{1+1/k})$.

Proof: Part (1) follows by the same argumentation as in Lemma 2.9 of [5]. Consider Part (2). Since $\mu = 1/k$, it holds that H_k obtained at the end of phase (S1) contains $O(n^{1+1/k})$ edges. We now bound the number of edges added in phase (S2), $|E_2| = \ell_t \cdot |Z_t| \cdot |Z_{k-1-t}|$, hence $|E_2| = \ell_t \cdot n^{1-\tau/k} \cdot n^{1-(k-1-\tau)/k} = \ell_t \cdot n^{1+1/k}$. Finally, in phase (S3), $(2k-1) \cdot |Z_\tau| \cdot |Z_{k-1-\tau}|$ edges are added for every $\tau \in \{0, \ldots, k-1\} \setminus \{t, t'\}$ and in addition, $\ell_t \cdot n^{1-t/k} \cdot n^{1-t'/k}$ edges connecting closest pairs in the clusters of $\mathcal{C}_t$ and $\mathcal{C}_{t'}$ are added as well. Since $|Z_\tau| = n^{1-\tau/k}$ and $|Z_{k-1-\tau}| = n^{1-(k-1-\tau)/k}$, it holds that $|E_3| = O(k^2 \cdot n^{1+1/k})$. ∎

We now turn to establish correctness. The following notation is useful in our analysis. A vertex is ℓ-clustered if it belongs to some $C \in \mathcal{C}_\ell$. For an ℓ-clustered vertex v, let $z_\ell(v) \in Z_\ell$ denote the cluster center of v in $\mathcal{C}_\ell$. An edge $e = (u, v)$ is ℓ-clustered if both u and v are ℓ-clustered, otherwise it is ℓ-unclustered. The next lemma (see Thm. 3.1 in [5]) plays a major role in our stretch analysis.

Lemma 2. *For every ℓ-unclustered edge (u, v), it holds that* $\mathrm{dist}(u, v, H) \leq 2\ell - 1$.

Proof: Let ℓ' be the minimum index such that either u or v was unclustered in $\mathcal{C}_{\ell'}$ (clearly, $\ell' \leq \ell$) and without loss of generality let the unclustered vertex be u. Then by Phase (T3), there is an edge (u, w) in H_ℓ from u to some vertex in $C_{\ell'-1}(v)$. In addition, there is a path in H_ℓ from w to v of length at most $2(\ell'-1)$, twice the radius of $C_{\ell'-1}(v)$. Since $\ell' \leq \ell$, it follows that $\mathrm{dist}(u, v, H) \leq 2\ell - 1$. ∎

We begin by considering the stretch between pairs of cluster centers $Z_{t'} \times Z_t$.

Lemma 3. *For every pair of cluster centers $z_i, z_j \in Z_{t'} \times Z_t$ it holds that*
(1) If $\mathrm{dist}(z_i, z_j, G) \leq \ell_t$, *then* $\mathrm{dist}(z_i, z_j, H) = \mathrm{dist}(z_i, z_j, G)$.
(2) If $\mathrm{dist}(z_i, z_j, G) > \ell_t$, *then* $\mathrm{dist}(z_i, z_j, H) \leq 2t \cdot (\mathrm{dist}(z_i, z_j, G) + 1) - \ell_t$.

Proof: Fix a cluster center $z_i \in Z_{t'}$ and let the cluster centers $Z_t = \{z_1, \ldots, z_\ell\}$ be ordered in nondecreasing distance from z_i, i.e., $\mathrm{dist}(z_i, z_1, G) \leq \mathrm{dist}(z_i, z_2, G) \leq \ldots \leq \mathrm{dist}(z_i, z_\ell, G)$. Assume, towards contradiction, that the the claims do not hold, and let z_j be the first center in the ordering for which one of the claims does not hold. Define $D_G = \mathrm{dist}(z_i, z_j, G)$ and $D_H = \mathrm{dist}(z_i, z_j, H)$. Since the last ℓ_t edges of the path $\pi(z_i, z_j)$ are taken into H, it holds that z_j does not satisfy claim (2) and hence

$$D_G = |\pi(z_i, z_j)| > \ell_t \ . \tag{5}$$

We now distinguish between cases depending on the type of the edges missing from $\pi(z_i, z_j)$ in H.
Case (a): all missing edges in $\pi(z_i, z_j) \setminus H$ are t-unclustered. By Lemma 2 and Eq. (5), we have that $D_H \leq (2t-1)D_G < 2t \cdot D_G - \ell_t$.
Next, consider the complementary case (b): some of the missing edges of $\pi(z_i, z_j) \setminus H$ are t-clustered. Let $e = (x_1, x_2)$ be the last missing edge on the path $\pi(z_i, z_j)$ (the edge closest to z_j) which is t-clustered, hence x_2 is t-clustered. Let $z' = z_t(x_2)$ be the cluster center of x_2. For an illustration, see Fig. 1. We first claim that z' precedes z_j in the ordering. To see this, recall that the last ℓ_t edges of the path $\pi(z_i, z_j)$ were taken into H , thus $\text{dist}(z_i, x_2, G) \leq \text{dist}(z_i, z_j, G) - \ell_t = D_G - \ell_t$. Since the radius of each cluster in $\mathcal{C}_t$ is at most t, we get that

$$\text{dist}(z_i, z', G) \leq \text{dist}(z_i, x_2, G) + t \leq D_G - \ell_t + t < D_G \ , \tag{6}$$

where the strict inequality follows by Eq. (4). This strict inequality implies that indeed z' precedes z_j in the ordering and by the definition of z_j it follows that z' satisfies the lemma. Consider the alternative $z_i - z_j$ path $P = P_1 \circ P_2 \circ P_3$ where $P_1 \in SP(z_i, z', H)$, $P_2 \in SP(z', x_2, H)$ and $P_3 \in SP(x_2, z_j, H)$. Since $P \subseteq H$, it remains to bound its length. Since z' is the cluster center of x_2, $|P_2| \leq t$. Since P_3 is free from missing t-clustered edges, by Lemma 2, $|P_3| \leq (2t-1) \cdot \text{dist}(x_2, z_j, G)$. To bound the length of P_1 we distinguish between two subcases depending on the length of $\pi(z_i, z')$. Subcase (b1): $|\pi(z_i, z')| \leq \ell_t$. Then, $\text{dist}(z_i, z', H) = \text{dist}(z_i, z', G)$ and overall we have that

$$\begin{aligned} D_H \leq |P| = |P_1| + |P_2| + |P_3| \ &\leq \ \text{dist}(z_i, z', G) + t + (2t-1)\text{dist}(x_2, z_j, G) \\ &\leq D_G - \ell_t + 2t + (2t-1) \cdot \text{dist}(x_2, z_j, G) \ \leq \ 2t \cdot (D_G + 1) - \ell_t \ , \end{aligned}$$

where the second inequality follows by the second inequality of Eq. (6). This contradicts the fact that z_j violates the claim.
Subcase (b2): $\text{dist}(z_i, z', G) \geq \ell_t + 1$. Since z' satisfies the lemma,

$$\text{dist}(z_i, z', H) \leq 2t(\text{dist}(z_i, z', G) + 1) - \ell_t \ . \tag{7}$$

Overall, we get that

$$\begin{aligned} D_H &\leq |P| = |P_1| + |P_2| + |P_3| \\ &\leq 2t(\text{dist}(z_i, z', G) + 1) - \ell_t + t + (2t-1) \cdot \text{dist}(x_2, z_j, G) && (8) \\ &\leq \text{dist}(z_i, z', G) + (2t-1) \cdot \text{dist}(z_i, z', G) + 3t - \ell_t + (2t-1) \cdot \text{dist}(x_2, z_j, G) \\ &\leq D_G - \ell_t + 4t + (2t-1) \cdot \text{dist}(z_i, z', G) - \ell_t + (2t-1) \cdot \text{dist}(x_2, z_j, G) && (9) \\ &\leq D_G - \ell_t + 4t + (2t-1) \cdot (\text{dist}(z_i, x_2, G) + t) - \ell_t + (2t-1) \cdot \text{dist}(x_2, z_j, G) \\ &\leq 2t \cdot D_G - \ell_t, && (10) \end{aligned}$$

where Eq. (8) follows by Eq. (7), Ineq. (9) follows by the inequality of Eq. (6) and the penultimate inequality follows by the fact that z' and x_2 belong to the same radius-t cluster. In contradiction to the definition of z_j. Claim (2) follows. ∎

Let V_i be the set of i-clustered vertices. We now turn to bound the stretch between pairs of vertices in $V_{t'} \times V_t$.

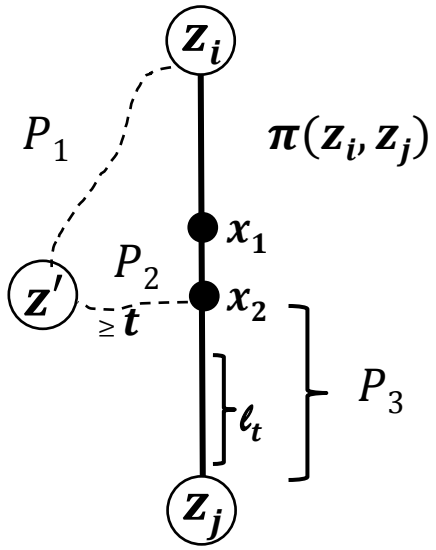

Fig. 1. Illustration of case (b) for Lemma 3(2). If not all edges on $\pi(z_i, z_j)$ are t-clustered, then the induction assumption for cluster centers preceding z_j in the ordering can be used; this is due to the fact that the last ℓ_t edges of $\pi(z_i, z_j)$ were taken into H. The dashed lines represent bypasses in H.

Lemma 4. *For every non-neighboring vertex pair $x_1 \in V_{t'}$ and $x_2 \in V_t$ (i.e., such that $\mathrm{dist}(x_1, x_2, G) \geq 2$), $\mathrm{dist}(x_1, x_2, H) \leq k \cdot \mathrm{dist}(x_1, x_2, G)$.*

Proof: We consider the following cases.
Case (1): $\mathrm{dist}(C_{t'}(x_1), C_t(x_2), G) \leq \ell_t$. Then again by phase (S3), there exists an $w - z$ path P in H such that $w \in C_{t'}(x_1)$ and $z \in C_t(x_2)$ and $|P| = \mathrm{dist}(w, z, G) \leq \mathrm{dist}(x_1, x_2, G)$. Hence

$$\begin{aligned}
\mathrm{dist}(x_1, x_2, H) &\leq \mathrm{dist}(x_1, w, H) + \mathrm{dist}(w, z, H) + \mathrm{dist}(z, x_2, H) \\
&= \mathrm{dist}(x_1, w, H) + \mathrm{dist}(w, z, G) + \mathrm{dist}(z, x_2, H) \\
&\leq 2t' + \mathrm{dist}(w, z, G) + 2t \leq 2(k-1) + \mathrm{dist}(x_1, x_2, G) \quad (11) \\
&\leq k \cdot \mathrm{dist}(x_1, x_2, H) \ ,
\end{aligned}$$

where Eq. (11) follows by the fact that x_1 (resp., x_2) and w (resp., z) belongs to the same cluster of diameter $2t'$ (resp., $2t$), and w and z are the closest pair in $C_{t'}(x_1)$ and $C_t(x_2)$. Finally, the last inequality follows as $\mathrm{dist}(x_1, x_2, H) \geq 2$.

Case (2): $\mathrm{dist}(C_{t'}(x_1), C_t(x_2), G) > \ell_t$. Let $z_1 = z_{t'}(x_1)$ (resp., $z_2 = z_t(x_2)$) be the cluster center of x_1 (resp., x_2) in the clustering $\mathcal{C}_{t'}$ (resp., $\mathcal{C}_t$). Since $\mathrm{dist}(z_1, z_2, G) \geq \mathrm{dist}(C_{t'}(x_1), C_t(x_2), G) > \ell_t$, by Lemma 3, it holds that $\mathrm{dist}(z_1, z_2, H) \leq 2t \cdot (\mathrm{dist}(z_1, z_2, G) + 1) - \ell_t$. Hence,

$$\begin{aligned}
\mathrm{dist}(x_1, x_2, H) &\leq \mathrm{dist}(x_1, z_1, H) + \mathrm{dist}(z_1, z_2, H) + \mathrm{dist}(z_2, x_2, H) \\
&\leq 2t + \mathrm{dist}(z_1, z_2, H) + 2t' && (12) \\
&\leq 2(k-1) + 2t \cdot \mathrm{dist}(z_1, z_2, G) + 2t - \ell_t && (13) \\
&\leq 2(k-1) + 2t \cdot (\mathrm{dist}(x_1, x_2, G) + 2(k-1)) - \ell_t && (14) \\
&\leq k \cdot \mathrm{dist}(x_1, x_2, G) \ ,
\end{aligned}$$

where Eq. (12) follows by the intercluster connections in clusters of diameter $2t$ and $2t'$, Eq. (13) follows by Eq. (3) and by Cl. 3, Eq. (14) follows by the triangle inequality using the fact that z_1 and x_1 are in the same cluster of diameter $2t'$ and z_2 and x_2 are in the same cluster of diameter $2t$. The claim follows. ∎

Finally, we are ready to to bound the distance for every $u, v \in V \times V$. We have the following.

Lemma 5. *For every $(u, v) \in V \times V$:*
(1) $\mathrm{dist}(u, v, H) \leq 2k - 1$ if $(u, v) \in E$.
(2) $\mathrm{dist}(u, v, H) \leq k \cdot \mathrm{dist}(u, v, G)$, otherwise.

Proof: Part (1) follows by the fact that H_k (the resulting subgraph of step (S1)) is a $(2k-1)$ spanner (see Thm. 3.1 of [5]). Consider part (2) and let $\pi(u, v)$ be the $u - v$ shortest path in G. Clearly, if u is t'-clustered and v is t-clustered or vice versa then the lemma follows by Lemma 4. Hence, it remains to consider the complementary case. From now on, assume without loss of generality that u is t_u-*unclustered* for a fixed $t_u \in \{t, t'\}$ and define $t_v = k - 1 - t_u$ (v can be either t_v-clustered or t_v-unclustered). If all the edges in $\pi(u, v)$ are t-unclustered then by Lemma 2, $\mathrm{dist}(u, v, H) \leq (2t-1) \cdot \mathrm{dist}(u, v, G)$ and we are done. Hence, from now on, assume that $\pi(u, v)$ contains at least two t-clustered vertices. In the same manner, if all the edges in $\pi(u, v)$ are t'-unclustered then by Lemma 2, $\mathrm{dist}(u, v, H) \leq (2t'-1) \cdot \mathrm{dist}(u, v, G)$ and we are done. Therefore $\pi(u, v)$ contains at least two t_u-clustered vertices and at least two t_v-clustered vertices. Let x_1 be the closest t_u-clustered vertex to u and let x_2 be the closest t_v-clustered vertex to v. Note that by definition, $u \neq x_1$. To prove Claim (2) of the lemma, we partition the path $\pi(u, v)$ into three segments according to x_1 and x_2. Let $\pi_1 = \pi(u, x_1)$, $\pi_2 = \pi(x_1, x_2)$ and $\pi_3 = \pi(x_2, v)$. By Lemma 2, and recalling that $t' \leq t$ it holds that

$$\begin{aligned}
\mathrm{dist}(u, x_1, H) \quad &\leq \quad (2t-1) \cdot \mathrm{dist}(u, x_1, G) \qquad (15) \\
&\text{and} \quad \mathrm{dist}(x_2, v, H) \leq (2t-1) \cdot \mathrm{dist}(x_2, v, G).
\end{aligned}$$

First consider the case where $\pi_1 \circ \pi_2 \subseteq \pi(u, v)$ (e.g., in such a case either $x_1 = x_2$ or x_1 is strictly closer to v than x_2). By Eq. (15), it holds that every edge $(y_1, y_2) \in \pi(u, v)$ is either t-unclustered or t'-unclustered and hence

$\mathrm{dist}(y_1, y_2, H) \leq 2t-1$, concluding that $\mathrm{dist}(u, v, H) \leq k \cdot \mathrm{dist}(u, v, G)$ as desired. From now on, we therefore consider the case where x_1 appears on $\pi(u, v)$ *before* x_2 (i.e., closer to u) and thus $u \neq x_1 \neq x_2$. We now consider the following case analysis.

Case (a) $\mathrm{dist}(x_1, x_2, G) = 1$ *and* $\mathrm{dist}(u, x_2, G) \leq 2k-1$. Let $j \in \{0, \ldots, k-1-t_v\}$ be the maximum number satisfying that x_2 is $(t_v + j)$-clustered. Since x_2 is t_v-clustered, there exists an j in this range. We now consider two subcases; for a schematic illustration see Fig. 2.

Case (a1): u is $(k-1-(t_v+j))$-clustered. Let $k' = (k-1-(t_v+j))$. Since $\mathrm{dist}(u, x_2, G) \leq 2k$, it also holds that $\mathrm{dist}(C_{k'}(u), C_{t_v+j}(x_2), G) \leq 2k$, hence by Phase (S3), there exists a $w-z$ path P in H where $w \in C_{k'}(u)$ and $z \in C_{t_v+j}(x_2)$ and $|P| = \mathrm{dist}(w, z, G) \leq \mathrm{dist}(u, x_2, G)$. Recall that as $|P| \leq 2k$, it was taken entirely into H. We therefore have the following.

$$
\begin{aligned}
\mathrm{dist}(u, v, H) &= \mathrm{dist}(u, x_2, H) + \mathrm{dist}(x_2, v, H) \\
&\leq \mathrm{dist}(u, w, H) + \mathrm{dist}(w, z, H) + \mathrm{dist}(z, x_2, H) \\
&\quad + (2t-1) \cdot \mathrm{dist}(x_2, v, G) && (16) \\
&\leq 2k' + \mathrm{dist}(w, z, G) + 2(t_v + j) + (2t-1) \cdot \mathrm{dist}(x_2, v, G) && (17) \\
&\leq 2k' + \mathrm{dist}(u, x_2, G) + 2(t_v + j) + (2t-1) \cdot \mathrm{dist}(x_2, v, G) && (18) \\
&\leq 2(k-1) + \mathrm{dist}(u, x_2, G) + (2t-1) \cdot \mathrm{dist}(x_2, v, G) && (19) \\
&\leq k \cdot \mathrm{dist}(u, x_2, G) + (2t-1) \cdot \mathrm{dist}(x_2, v, G) && (20) \\
&\leq k \cdot \mathrm{dist}(u, v, G)\ ,
\end{aligned}
$$

where Eq. (16) follows by Eq. (15), Eq. (17) follows by Lemma 2 and by the fact that u is k'-clustered and x_2 is $t_v + j$ clustered, Eq. (18) follows by the fact that w and z are the closest pair in $C_{k'}(u)$ and $C_{t_v+j}(x_2)$, finally the penultimate inequality holds as $\mathrm{dist}(u, x_2, G) \geq 2$ (since $u \neq x_1 \neq x_2$).

Case (a2): u is $(k-1-(t_v+j))$-unclustered. Let $k' = k-1-(t_v+j)$. Since all vertices are 0-clustered, it holds that $k' \geq 1$ and hence $j \leq k - t_v - 2$. By the definition of j, we have that x_2 is $(t_v + j + 1)$-unclustered (as $j \leq k - t_v - 2$, it holds that $(t_v + j + 1) \leq k-1$). Also, recall that $x_1 \neq u$. Let $\pi(u, v) = [u = u_0, u_1, \ldots, u_q = v]$. Since u_0 is k'-unclustered, by Lemma 2, it holds $\mathrm{dist}(u_0, u_1, H) \leq 2k' - 1$. In the same manner, since x_2 is $(t_v + j + 1)$-unclustered, it holds that $\mathrm{dist}(x_1, x_2, H) \leq 2(t_v + j + 1) - 1$. Putting all together, we have that

$$
\begin{aligned}
\mathrm{dist}(u, v, H) &= \mathrm{dist}(u_0, u_1, H) + \mathrm{dist}(u_1, x_1, H) + \mathrm{dist}(x_1, x_2, H) + \mathrm{dist}(x_2, u_q, H) \\
&\leq 2k' - 1 + (2t-1) \cdot \mathrm{dist}(u_1, x_1, G) + 2(t_v + j + 1) - 1 && (21) \\
&\quad + (2t-1) \cdot \mathrm{dist}(x_2, u_q, G) \\
&= 2(k-1) + (2t-1) \cdot (\mathrm{dist}(u_1, x_1, G) + \mathrm{dist}(x_2, u_q, G)) \\
&= (k-1) \cdot (\mathrm{dist}(u_0, u_1, G) + \mathrm{dist}(x_1, x_2, G)) && (22) \\
&\quad + (2t-1) \cdot (\mathrm{dist}(u_1, x_1, G) + \mathrm{dist}(x_2, u_q, G) \leq k \cdot \mathrm{dist}(u, v, G)\ ,
\end{aligned}
$$

where Eq. (21) follows by using Lemma 2 for the k'-unclustered edge (u_0, u_1) and for the $(t_v + j + 1)$-unclustered edge (x_1, x_2), and by plugging Eq. (15). Note that it might be the case that $x_1 = u_1$ but by the current case it holds that $u_0 \neq x_1 \neq x_2$.

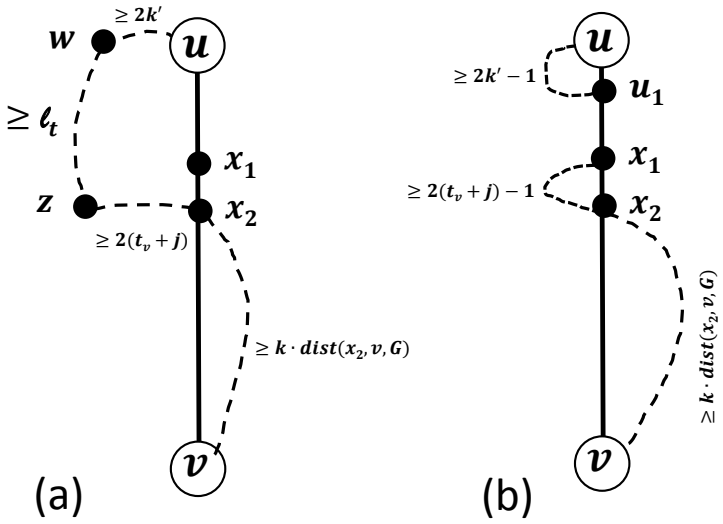

Fig. 2. Illustration of Case (a) for Lemma 5. The dashed lines represent bypasses in H. (a) Case (a1). Since u is k'-clustered and x_2 is $(t_v + j)$-clustered and in addition, $\text{dist}(u, v, G) \leq \ell_t$ it holds that the shortest path between the corresponding clusters has been added to H. (b) Case (a2). Since u is k'-unclustered, its immediate edges, specifically, (u_0, u_1) has small stretch in H that compensates for the extra large stretch of the edge (x_1, x_2) in H.

Case (b): $\text{dist}(x_1, x_2, G) = 1$ *and* $\text{dist}(u, x_2, G) \geq 2k$. Using Part (1) of this claim, we have that

$$\begin{aligned}
\text{dist}(u, v, H) &= \text{dist}(u, x_1, H) + \text{dist}(x_1, x_2, H) + \text{dist}(x_2, v, H) && (23)\\
&\leq (2t-1) \cdot \text{dist}(u, x_1, G) + \text{dist}(x_1, x_2, H) \\
&+ (2t-1) \cdot \text{dist}(x_2, v, G) && (24)\\
&\leq (2t-1) \cdot \text{dist}(u, x_1, G) + 2k - 1 \\
&+ (2t-1) \cdot \text{dist}(x_2, v, G) && (25)\\
&\leq (2t-1) \cdot \text{dist}(u, v, G) + \text{dist}(u, v, G) \leq 2t \cdot \text{dist}(u, v, G)\ . && (26)
\end{aligned}$$

where Eq. (24) follows by Eq. (15), Eq. (25) follows by Part (1) of this claim and Eq. (26) follows by the fact that $\mathrm{dist}(u,v,G) \geq \mathrm{dist}(u,x_2,G) \geq 2k-1$. The claim holds.

Case (c): $\mathrm{dist}(x_1,x_2,G) \geq 2$. This case follows by using Lemma 4(b) and Eq. (15). The claim follows. ∎

Running time. Begin with step (S1). By Theorem 3.1 of [5], Algorithm $\mathsf{Cluster}(G,k,\mu)$ can be implemented in $O(k \cdot m)$ deterministic time. Step (S2) can be implemented in linear time (in the size of the output). Finally, step (S3) requires computing the shortest-path distances between cluster pairs. This can be done by computing the all-pairs-shortest-paths in $O(n^\omega)$, where $\omega < 2.373$ denotes the matrix multiplication exponent. This completes the proof of Thm. 1.

3 Sourcewise spanners

In this section, we provide several constructions for sourcewise spanners and emulators.

3.1 Upper bound for multiplicative stretch

For clarity of presentation, we describe a randomized construction whose output spanner has $O(k^2 \cdot n^{1+\varepsilon/k})$ edges in expectation. Using [5], this construction can be derandomized with the same bound on the number of edges. We begin by considering Thm. 2 and show the construction of $(2k-1,S)$ sourcewise spanner which enjoys a "hybrid" stretch, though in a weaker sense than in Sec. 2. Specifically, we show that the neighbors of S enjoy a multiplicative stretch $2k-1$ and the remaining pairs enjoy a multiplicative stretch of $2k-2$.

The algorithm The first phase of Algorithm $\mathsf{ConsSWSpanner}$ applies Algorithm $\mathsf{Cluster}(G,k,\mu)$ for $\mu = \varepsilon/k$, resulting in a sequence of $k+1$ clusterings $\mathcal{C}_0,\ldots,\mathcal{C}_k$ and a cluster graph $H_k \subseteq G$. In the second phase of the algorithm, it considers the collection of $S \times Z_{k-1}$ shortest paths $\mathcal{P} = \{\pi(s_j,z_i) \mid s_j \in S \text{ and } z_i \in Z_{k-1}\}$. Starting with $H = H_k$, for each path $\pi(s_j,z_i) \in \mathcal{P}$, it adds to H the ℓ_k last edges of $\pi(s_j,z_i)$ (closest to z_i). Set

$$\ell_k = 2k^2 + 3k \quad \text{and} \quad \mu = \varepsilon/k \ . \tag{27}$$

Analysis. Hereafter, we mainly focus on the clustering $\mathcal{C}_{k-1}$ and denote a vertex v as *clustered* iff $v \in V_{k-1}$, i.e., if v belongs to the clustering $\mathcal{C}_{k-1}$. Recall that each of the $m_{k-1} = n^{1-(k-1)\cdot\mu}$ clusters $C_i \in \mathcal{C}_{k-1}$ is centered at some vertex $z_i \in Z_{k-1}$, and in addition, every vertex v satisfying that $\mathrm{dist}(v,Z_{k-1},G) \leq k-1$ is contained in exactly one cluster $C_i \in \mathcal{C}_{k-1}$ such that

$$\mathrm{dist}(v,z_i,G) \leq k-1 \ . \tag{28}$$

An edge $e=(x,y)$ is said to be *clustered* iff both its endpoints x and y are clustered, i.e., $x,y \in V_{k-1}$. By Lemma 1(1), we have the following.

Observation 7 $\mathbb{E}(|H|) = O(k^2 \cdot n^{1+\varepsilon/k})$.

Next, we turn to consider correctness. Call a subgraph $H' \subseteq G$ *happy* iff every missing edge $e = (x, y) \in E(G) \setminus E(H')$ is either clustered or has stretch at most $2(k-1)-1$ in H, i.e., $\text{dist}(x, y, H) \leq 2k-3$. By Lemma 2, we have the following.

Observation 8 *H_k is happy (i.e., for every edge $e = (x, y) \in G \setminus H_k$, either e is clustered or $\text{dist}(x, y, H_k) \leq 2k - 3$).*

We now bound the stretch on $S \times Z_{k-1}$, i.e., between sources and cluster centers.

Lemma 6. *For every $s_j \in S$ and $z_i \in Z_{k-1}$:*
If $\text{dist}(s_j, z_i, G) \leq \ell_k$ then $\text{dist}(s_j, z_i, H) = \text{dist}(s_j, z_i, G)$, else $\text{dist}(s_j, z_i, H) \leq 2(k-1) \cdot (\text{dist}(s_j, z_i, G) + 1) - \ell_k$.

Proof: Fix $s_j \in S$ and let the cluster centers $Z_k = \{z_1, \ldots, z_m\}$ be ordered in *nondecreasing* distance from s_j, i.e., $\text{dist}(s_j, z_1, G) \leq \text{dist}(s_j, z_2, G) \leq \ldots \leq \text{dist}(s_j, z_m, G)$. (Note that the ordering of the centers plays a role only in the analysis and not in the algorithm itself.) Assume, towards contradiction, that the claim does not hold, and let $\widehat{z}$ be the first center in the ordering for which the claim does not hold. We first claim that

$$D_G = \text{dist}(s_j, \widehat{z}, G) = |\pi(s_j, \widehat{z})| > \ell_k \ , \tag{29}$$

since by construction, if $|\pi(s_j, \widehat{z})| \leq \ell_k$ then $\pi(s_j, \widehat{z})$ is taken entirely into H and $\text{dist}(s_j, \widehat{z}, H) = \text{dist}(s_j, \widehat{z}, G)$, contradiction. The contradiction assumption is therefore that $D_H = \text{dist}(s_j, \widehat{z}, H)$ satisfies $D_H > 2(k-1) \cdot (D_G + 1) - \ell_k$. We now distinguish between two cases depending on the type of missing edges in $\pi(s_j, \widehat{z}) \setminus H$ (by the contradictory assumption such missing edges exist). Case (a) is where none of the missing edges is clustered. By Obs. 8, we have that $D_H \leq (2k-3) \cdot D_G \leq 2(k-1) \cdot D_G - \ell_k$, where the last inequality follows by Eq. (29).

Next, consider the complementary case (b) where some of the missing edges of $\pi(s_j, \widehat{z}) \setminus H$ are clustered. Let $e = (x_1, x_2)$ be the last missing edge on the path $\pi(s_j, \widehat{z})$ (the edge closest to $\widehat{z}$) which is clustered, hence $x_2 \in V_{k-1}$. Consider the cluster center z' of the cluster C of x_2. We claim that z' precedes $\widehat{z}$ in the ordering. To see this, recall that the last ℓ_k edges of the path $\pi(s_j, \widehat{z})$ were taken into H, thus $\text{dist}(s_j, x_2, G) \leq \text{dist}(s_j, \widehat{z}, G) - \ell_k = D_G - \ell_k$. Combining with Eq. (28),

$$\text{dist}(s_j, z', G) \leq \text{dist}(s_j, x_2, G) + k - 1 \leq D_G - \ell_k + k - 1 < D_G \ , \tag{30}$$

where the first inequality follows by Eq. (28) and the strict inequality follows by Eq. (27). This strict inequality implies that indeed z' precedes $\widehat{z}$ in the ordering and by the definition of $\widehat{z}$ (i.e., the first center that violates the claim) it follows that z' satisfies the lemma.

Consider the following alternative $s_j - \widehat{z}$ path $P = P_1 \circ P_2 \circ P_3$ which consists of three segments: a source-center path $P_1 \in SP(s_j, z', H)$, an intra-cluster path $P_2 \in SP(z', x_2, H)$ and $P_3 \in SP(x_2, \widehat{z}, H)$. Since $P \subseteq H$ by definition, it remains to bound its length. By Eq. (28), $|P_2| \leq k-1$. By Obs. 8, since P_3 is free from missing clustered edges, $|P_3| \leq (2k-1) \cdot \mathrm{dist}(x_2, \widehat{z}, G)$. To bound the length of P_1 we distinguish between two cases depending on the size of $\pi(s_j, z')$. Case (a): $|\pi(s_j, z')| \leq \ell_k$. Then, $\mathrm{dist}(s_k, z', H) = \mathrm{dist}(s_k, z', G)$ and overall we have that

$$\begin{aligned} D_H &\leq |P| = |P_1| + |P_2| + |P_3| \\ &\leq \mathrm{dist}(s_j, z', G) + k - 1 + (2k-3) \cdot \mathrm{dist}(x_2, \widehat{z}, G) \\ &\leq D_G - \ell_k + 2(k-1) + (2k-3) \cdot \mathrm{dist}(x_2, \widehat{z}, G) \\ &\leq 2(k-1) \cdot (D_G + 1) - \ell_k \ , \end{aligned}$$

where the third inequality follows by Eq. (30). This contradicts the assumption that $\widehat{z}$ violates the claim. Case (b): $\mathrm{dist}(s_j, z', G) \geq \ell_k + 1$. Since z' satisfies the lemma,

$$\mathrm{dist}(s_j, z', H) \leq 2(k-1) \cdot (\mathrm{dist}(s_j, z', G) + 1) - \ell_k. \tag{31}$$

Overall, we get that

$$\begin{aligned} D_H &\leq |P| = |P_1| + |P_2| + |P_3| && (32) \\ &\leq 2(k-1) \cdot (\mathrm{dist}(s_j, z', G) + 1) - \ell_k \\ &\quad + k - 1 + (2k-3) \cdot \mathrm{dist}(x_2, \widehat{z}, G) && (33) \\ &= \mathrm{dist}(s_j, z', G) + (2k-3) \cdot \mathrm{dist}(s_j, z', G) + 3(k-1) - \ell_k \\ &\quad + (2k-3) \cdot \mathrm{dist}(x_2, \widehat{z}, G) \\ &\leq D_G - \ell_k + k - 1 + (2k-3) \cdot \mathrm{dist}(s_j, z', G) \\ &\quad + 3(k-1) - \ell_k + (2k-3) \cdot \mathrm{dist}(x_2, \widehat{z}, G) && (34) \\ &\leq D_G - \ell_k + k - 1 + (2k-3) \cdot (\mathrm{dist}(s_j, x_2, G) + k - 1) && (35) \\ &\quad + 3(k-1) - \ell_k + (2k-3) \cdot \mathrm{dist}(x_2, \widehat{z}, G) \\ &= 2(k-1) \cdot D_G - \ell_k \ , \end{aligned}$$

where Eq. (33) follows by Eq. (31), Eq. (34) follows by Eq. (30), and Eq. (35) follows by the fact that x_2 and z' are in the same cluster of diameter $k-1$ (see Eq. (28)). This is again in contradiction to the definition of $\widehat{z}$. The lemma follows. ∎

We are now ready to bound the $S \times V$ stretch also for vertices $V \setminus Z_{k-1}$. We have the following.

Lemma 7. *For every $(s_j, v) \in S \cdot V$, $\mathrm{dist}(s_j, v, H) \leq k' \cdot \mathrm{dist}(s_j, v, G)$ where $k' = 2k-1$ if $(s_j, v) \in E(G)$ and $k' = 2k-2$ otherwise.*

Proof: Fix $s_j \in S$ and assume, towards contradiction, that there exists some vertex $v \in V$ for which the claim does not hold. We distinguish between two cases. Case (a): all the missing edges on $\pi(s_j, v) \setminus H$ are unclustered. Then, by

Obs. 8, it follows that $\text{dist}(s_j, v, H) \leq (2k-3) \cdot \text{dist}(s_j, v, G)$. Case (b): there exists a missing edge which is clustered. Let $e = (x_1, x_2)$ be the last such missing edge on $\pi(s_j, v)$ (closest to v) and let z' be the cluster center of the cluster C of x_2. Consider an $s_j - v$ path $P = P_1 \circ P_2 \circ P_3$ in H which consists of the following three segments: $P_1 \in SP(s_j, z', H)$, $P_2 \in SP(z', x_2, H)$ and $P_3 \in SP(x_2, v, H)$. Note that by Eq. (28) and Obs. 8,

$$|P_2| \leq k-1 \quad \text{and} \quad |P_3| \leq (2k-3) \cdot \text{dist}(x_2, v, G) \ . \tag{36}$$

Bounding the size of P_1 is more involved and requires the following case analysis which depends on $\text{dist}(s_j, z', G)$.

Case a: $\text{dist}(s_j, z', G) \leq \ell_k$. In this case, the path $\pi(s_j, z')$ was added in its entirety to H and hence $\text{dist}(s_j, z', H) = \text{dist}(s_j, z', G)$. Combining with Eq. (36) we have that

$$\begin{aligned} \text{dist}(s_j, v, H) \leq |P| &= |P_1| + |P_2| + |P_3| \\ &\leq \text{dist}(s_j, z', G) + k - 1 + (2k-3) \cdot \text{dist}(x_2, v, G) \\ &\leq \text{dist}(s_j, x_2, G) + (2k-3) \cdot \text{dist}(x_2, v, G) + 2(k-1) \ . \end{aligned} \tag{37}$$

We now look at two sub-cases. Sub-case (a1): $\text{dist}(s_j, x_2, G) \geq 2$, then $\text{dist}(s_j, x_2, G) + 2(k-1) \leq 2(k-1) \cdot \text{dist}(s_j, x_2, G)$ and hence by Eq. (37) we conclude that $\text{dist}(s_j, v, H) \leq 2(k-1) \cdot \text{dist}(s_j, v, G)$ as required. The complementary sub-case (a2) is where $\text{dist}(s_j, x_2, G) = 1$. (By definition, $e = (x_1, x_2)$ is missing and hence $s_j \neq x_2$ or $\text{dist}(s_j, x_2, G) > 0$). This sub-case now splits further. Case (a2i): $x_2 \neq v$, i.e., $\text{dist}(x_2, v, G) \geq 1$, then by Eq. (37), $|P| \leq 2(k-1) + 1 + (2(k-1)-1) \cdot \text{dist}(x_2, v, G) \leq 2(k-1) \cdot \text{dist}(s_j, v, G)$. Case (a2ii): $x_2 = v$ (i.e., $\text{dist}(x_2, v, G) = 0$). Then $|P| \leq \text{dist}(s_j, x_2, G) + 2(k-1)$, in contradiction to the fact the v violates the claim.

Case b: $\text{dist}(s_j, z', G) > \ell_k$. By Lem. 6 we have that $|P_1| \leq 2(k-1) \cdot (\text{dist}(s_j, z', G) + 1) - \ell_k$. Overall, by plugging Eq. (36), we have that

$$\begin{aligned} \text{dist}(s_j, v, H) \leq |P| &= |P_1| + |P_2| + |P_3| \\ &\leq 2(k-1) \cdot (\text{dist}(s_j, z', G) + 1) - \ell_k + k - 1 \\ &\quad + (2k-3) \cdot \text{dist}(x_2, v, G) \\ &\leq 2(k-1) \cdot (\text{dist}(s_j, x_2, G) + k - 1) + 3(k-1) - \ell_k \\ &\quad + (2k-3) \cdot \text{dist}(x_2, v, G) \\ &\leq 2(k-1)(\text{dist}(s_j, x_2, G) + \text{dist}(x_2, v, G)) \\ &= 2(k-1) \cdot \text{dist}(s_j, v, G) \end{aligned}$$

where the penultimate inequality follows by Eq. (28) and the last inequality follows by plugging Eq. (27), in contradiction to the definition of v. The lemma follows. ∎

We next analyze the running time and establish Thm. 5.

Running time. The operation of the algorithm is composed of k iterations. In each iteration τ, a cluster graph is constructed in $O(m)$ time (as in [5]). In the second phase of the algorithm, the $s_i - z_j$ paths of length at most ℓ_k edges are added to the spanner. This can viewed as constructing cluster graph of radius ℓ_k, but in this case, the clusters may not be disjoint (i.e., a vertex v belongs to the cluster of z_j if $\text{dist}(v, z_j, G) \leq \ell_k$). This can be implemented in $O(\ell_k \cdot m)$ time.

3.2 Lower bound for additive sourcewise spanners and emulators

We now turn to consider the lower bound side where we generalize the lower bound construction for additive spanners by Woodruff [33] to the sourcewise setting. In particular, we parameterize our bound for the $S \times V$ spanner in terms of the cardinality of the source set S. The basic idea underlying Woodruff's construction is to form a dense graph G by gluing (carefully) together many small complete bipartite graphs. For an additive stretch $2k-1 \geq 1$, the lower bound graph G consists of $k+1$ vertex levels, each with $O(n/k)$ vertices and $\Omega(n^{1+1/k})$ edges connecting the vertices of every two adjacent levels. In particular this is obtained by representing each vertex of level i as a coordinate in $\mathbb{Z}^{k+1}$, namely, $v = (a_1, \ldots, a_k, a_{k+1})$ and $a_j \in [1, O(n^{1/k})]$. Woodruff showed that if one omits in an additive spanner $H \subseteq G$, an $O(1/k)$ fraction of G edges, then there exists an $x - y$ path P in G of length k (i.e., x is on the first level and y is on the last level) whose all edges are omitted in H, and any alternative $x - y$ path in H is "much" longer than P. To adapt this construction to the sourcewise setting, some *asymmetry* in the structure of the $k+1$ levels should be introduced. In the following construction, the vertices of the first level correspond to the source set S, hence this level consists of $O(n^\varepsilon)$ vertices, while the remaining levels are of size $O(n/k)$. This is achieved by breaking the symmetry between the first coordinate a_1 and the remaining $k-1$ coordinates of each vertex $v = (a_1, \ldots, a_k, a_{k+1})$. Indeed, this careful minor adaptation in the graph definition is sufficient to generalize the bound, the analysis follows (almost) the exact same line as that of [33]. We show the following.

Theorem 9. *Let $1 \leq k \leq O(\ln r / \ln \ln r)$ for some integer $r \geq 1$. For every $\varepsilon \in [0,1]$, there exists an unweighted undirected graph $G = (V, E)$ with $|V| = \Theta(r^\varepsilon + kr)$ vertices and a source set $S \subseteq V$ of size $\Theta(r^\varepsilon)$ such that any $(2k-1, S)$-additive sourcewise spanner $H \subseteq G$ has $\Omega(r^{1+\frac{\varepsilon}{k}})$ edges. Similar bounds (up to factor $O(k)$) are achieved for $(2k-1, S)$-additive sourcewise emulators.*

Note that Thm. 9 implies Thm. 3, since $n = \Theta(r^\varepsilon + kr)$ and hence $r^{1+\frac{\varepsilon}{k}} = \Omega(n^{1+\frac{\varepsilon}{k}}/k)$. Note that by setting $\varepsilon = 1$, we get the exact same bounds as in Woodruff's construction.

The construction. Let $N_1 = \lceil r^{\varepsilon/k} \rceil$ and $N_2 = \lceil (r/N_1^{k-1}) \rceil$. The graph G consists of vertices composed of $k+1$ vertex-levels and connected through a series of k bipartite graphs. Each vertex $v = (a_1, a_2, \ldots, a_k, a_{k+1})$ represents a coordinate in $\mathbb{Z}^{k+1}$ where $a_{k+1} \in \{1, \ldots, k+1\}$ is the *level* of v. The range

of the other coordinates is as follows. For every $1 \leq j \leq k$, $a_j \in R_j$, where $R_1 = \{1, \ldots, N_1\}$ if $a_{k+1} = 1$ and $R_1 = \{1, \ldots, N_2\}$ otherwise. For $j \geq 2$, $R_j = \{1, \ldots, N_1\}$.
Edges in G join every level-i vertex $(a_1, \ldots, a_{i-1}, a_i, a_{i+1}, \ldots, a_k, i)$ to each of the level-$(i+1)$ vertices of the form $(a_1, \ldots, a_{i-1}, c, a_{i+1}, \ldots, a_k, i+1)$ for every $c \in \{1, \ldots, N_2\}$ if $i = 1$ and $c \in \{1, \ldots, N_1\}$ for $i \geq 2$. Let $L_i = \{(a_1, \ldots, a_k, i) \mid a_j \in R_j \text{ for } 1 \leq j \leq k\}$ be the set of vertices on the ith level and let $n_i = |L_i|$ denote their cardinality. Then since $k = O(\ln r / \ln \ln r)$ it holds that $n_1 = N_1^k \leq (r^{\frac{\varepsilon}{k}} + 1)^k \leq e^{(k+1)/(r^{\varepsilon/k})} = \Theta(r^\varepsilon)$. and for every $i \in \{2, \ldots, k+1\}$,

$$\begin{aligned} n_i = N_2 \cdot N_1^{k-1} &\leq (r/N_1^k + 1)(r^{\frac{\varepsilon}{k}} + 1)^{k-1} \leq 2r^{1-\varepsilon/k} \cdot e^{k/(r^{\varepsilon/k})} \\ &= r^{1-\varepsilon/k} \cdot \Theta(r^{\varepsilon/k}) = \Theta(r) \ , \end{aligned}$$

Overall, the total number of vertices is $|V(G)| = n_1 + k \cdot n_2 = \Theta(r^\varepsilon + k \cdot r)$.

Let g_i be the number of edges connecting the vertices of L_i to the vertices of L_{i+1}. Then $g_1 = N_2 \cdot n_1$ and $g_i = N_1 \cdot n_i$ for every $i \in \{2, \ldots, k\}$, thus $g_1 = g_2 = \ldots = g_k$. Hence $|E(G)| = \sum_{i=1}^{k+1} g_i = k \cdot N_1^k \cdot N_2 = \Theta(k \cdot r^{1+\varepsilon/k})$. Let the source set S be the vertex set of the first level, i.e., $S = L_1$, hence $|S| = n_1 = \Theta(r^\varepsilon)$. The analysis of the construction follows the same line as that of [33]. For completeness, we briefly describe it. We begin by exploring the following *distortion property* on some spanner $\subseteq G$. Consider any subgraph H with fewer than $E(G)/k = N_1^k \cdot N_2 = \Theta(r^{1+\varepsilon/k})$ edges. We now show that there exist vertices $u \in L_1$ and $v \in L_{k+1}$ such that $\text{dist}(u, v, H) \geq \text{dist}(u, v, G) + 2k$. By Lemma 5 of [33], we have the following.

Lemma 8. *There exist $k+1$ vertices $v_1, \ldots, v_{k+1}$ such that for each i, $v_i \in L_i$ and the edge (v_i, v_{i+1}) is missing in H.*

Choose $v_1, \ldots, v_k$ as in Lemma 8. Then $\text{dist}(v_1, v_k, G) \leq k+1$ since $v_1, \ldots, v_{k+1}$ is a path in G. The following lemma shows that $\text{dist}(v_1, v_{k+1}, H)$ is large.

Lemma 9 (Lemma 6 of [33]). *Any path in G from v_1 to v_{k+1} of length less than $3k$ contains an edge (v_i, v_{i+1}) for some $i \in \{1, \ldots, k\}$, and further this is the only path edge on the ith level.*

Proof: Let P be any path from v_1 to v_{k+1} in G. Let i be any level $i \in \{1, \ldots, k+1\}$. Starting from v_1, traverse the edges of P one by one. After encountering an even number of edges in level i as we walk along P, we must be in level j such that $j \leq i$. Thus as P starts with a level-1vertex and ends with a level- $(k+1)$ vertex, P must contain an odd number of edges in each level i.

Therefore, if the length of P is less than $3k$, then by the pigeonhole principle there is an i for which P contains exactly one edge in level i (since otherwise P contains at least 3 edges in each level so $|P| \geq 3(k+1)$). Let (a, b) denote this edge. We claim that $(a, b) = (v_i, v_{i+1})$. To see this, first note that the last $k - (i-1)$ coordinates (not including the level coordinate) of a must agree with those of v_1 since (i) P begins at v_1, (ii) all edges in P preceding (a, b) are in level $j < i$, and (iii) an edge in level j, for any $1 \leq j \leq k$, may only modify

the jth coordinate of its endpoints. Moreover, as (a, b) is the only edge in level i, P cannot return to any level $j < i$. Therefore, since P ends at v_{k+1}, the first $i-1$ coordinates of a must agree with those of v_{k+1}. By definition then, $a = v_i$. As only edges on the ith level affect the ith coordinate, necessarily $b = v_{i+1}$, as otherwise another edge in level i would be needed to correct the ith coordinate so that P could reach v_k. This proves the lemma. ∎

Corollary 1. $\text{dist}(v_1, v_{k+1}, H) \geq 3k$ *hence* $\text{dist}(v_1, v_{k+1}, H) > \text{dist}(v_1, v_{k+1}, G) + 2k - 1$.

Proof: By the previous lemma, any path in G of length less than $3k$ contains an edge of the form (v_i, v_{i+1}), and by our choice of $v_1, \ldots, v_k$, this edge is missing in H so the path does not occur in H. ∎

It follows that any subgraph of G with less than $|E(G)|/(k+1) = \Omega(r^{1+\varepsilon/k})$ edges distorts some $S \times V$ pair by at least an additive $2k$, so it is not a $(2k-1, S)$-additive sourcewise spanner.

Theorem 10. *Let $1 \leq k \leq O(\ln r / \ln \ln r)$ be an integer. For every $\varepsilon \in [0, 1]$, there exists an unweighted undirected graph $G = (V, E)$ with $|V| = \Theta(r^\varepsilon + kr)$ vertices and a source set $S \subseteq V$ of size $\Theta(r^\varepsilon)$ such that any $(2(k-1), S)$-additive emulator $H = (V, F)$ has $\Omega(1/k \cdot r^{1+\varepsilon/k})$ edges.*

The graph G that achieves the lower bound is the same as in the proof of Thm. 9. The correctness of the construction follows immediately from the proof of Thm. 8 in [33]. In short, one can show that the diameter of G is at most $O(k)$. In addition, note that for any edge $(u, v) \in E(H)$ it is never optimal for the weight of (u, v) to be larger than $\text{dist}(u, v, G)$. Combining these two observation implies that any weighted edge (u, v) in the emulator H can be replaced by at most $O(k)$ edges in G, corresponding to the $u - v$ shortest-path in G. This results in a *subgraph* H' of G which according to Thm. 9 is of size $\Omega(1/k \cdot r^{1+\varepsilon/k})$. It follows that the emulator H is sparser by a factor of at most $O(k)$, i.e., $|E(H)| = \Omega(1/k^2 \cdot r^{1+\varepsilon/k})$, as required.

3.3 Upper bound for additive sourcewise spanners and emulators

Additive sourcewise emulators. Recall that an emulator $H = (V, F)$ for graph G is a (possibly) weighted graph induced on the vertices of G, whose edges are not necessarily contained in G. In Thm. 9, we showed that every $(2, S)$-additive sourcewise emulator for a subset $S \subseteq V$ has $\Omega(n^{1+\varepsilon/2})$ edges, where $|S| = O(n^\varepsilon)$. We now show that this is essentially tight (up to constants).

Theorem 11. *For every unweighted n-vertex graph $G = (V, E)$ and every subset $S \subseteq V$, there exists a (polynomially constructible) $(2, S)$-additive sourcewise emulator H of size $O(n^{1+\varepsilon/2})$ where $\varepsilon = \log |S| / \log n$.*

The following Fact from [12] is useful in the subsequent constructions.

Fact 12 ([12]) *There is a polynomial time algorithm* $\mathsf{Cluster}(G,\gamma)$ *that given a parameter* $\gamma \in [0,1]$ *and a graph* $G=(V,E)$ *constructs a collection of clusters* $\mathcal{C}$ *with at most* $n^{1-\gamma}$ *vertex-disjoint clusters, each of size* n^{γ}*, and a subgraph* G_C *of* G *with* $O(n^{1+\gamma})$ *edges such that (1) for any missing edge* $(u,v) \in E(G) \setminus E(G_C)$*,* u *and* v *belong to two different clusters and (2) the diameter of each cluster (i.e., the distance in* G_C *between any two vertices of the cluster) is at most* 2*.*

Proof of Theorem 11. The weighted graph emulator $H=(V,F)$ is constructed as follows. Let $W : F \rightarrow \mathbb{R}$ be the weights of H edges. Apply Fact 12 to construct a clustering $\mathcal{C}$ and the clustering graph $G_C = \mathsf{Cluster}(G,\gamma)$ for $\gamma = \varepsilon/2$. For every source $s_i \in S$ and cluster $C_j \in \mathcal{C}$, define $z_{i,j} \in C_j$ to be the closest vertex in cluster C_j to the source s_i, i.e., $z_{i,j} \in C_j$ satisfies that $\mathrm{dist}(s_i, z_{i,j}, G) = \mathrm{dist}(s_i, C_j, G)$, and add to H an edge $e_{i,j}$ between s_i and $z_{i,j}$ of weight $\mathrm{dist}(s_i, z_{i,j}, G)$. This concludes the construction of H. We first bound the size of H. By Fact 12, G_C has $O(n^{1+\varepsilon/2})$ edges and there are $|\mathcal{C}| = O(n^{1-\varepsilon/2})$ clusters. Hence, overall H consists of $|E(G_C)| = O(n^{1+\varepsilon/2})$ edges of weight 1 and $|S|\cdot|\mathcal{C}| = O(n^{1+\varepsilon/2})$ weighted edges as required. We now analyze the stretch. Consider some $s_i - v_j$ shortest path $\pi(s_i, v_j)$. If $\pi(s_i, v_j) \subseteq G_C$, then $\mathrm{dist}(s_i, v_j, H) = \mathrm{dist}(s_i, v_j, G)$ and the claim holds. Else, let $e=(w_1,w_2)$ be last missing edge on $\pi(s_i, v_j) \setminus G_C$, i.e., the edge closest to v_j that does not appear in the clustering graph G_C. By Fact 12, w_2 is clustered. Let $C = C(w_2)$ be the cluster of w_2 and let $z \in C$ be the closest vertex to s_i in the cluster C. By construction, H contains an edge $e' = (s_i, z)$ of weight $W(e') = \mathrm{dist}(s_i, z, G) = \mathrm{dist}(s_i, C, G)$. Let $P_1 = e'$, $P_2 \in SP(z, w_2, G_C)$, $P_3 = \pi(s_i, v_j)[w_2, v_j]$ and define the $s_i - v_j$ path $P = P_1 \circ P_2 \circ P_3$. By construction $P \subseteq H$. Let $W(P) = \sum_{e \in P} W(e)$ be weight length of P. Note that by Fc. 12, $W(P_2) \leq 2$. We then have the following.

$$\begin{aligned}\mathrm{dist}(s_i, v_j, H) &\leq W(P) = W(P_1) + W(P_2) + W(P_3) \leq \mathrm{dist}(s_i, C, G) + 2 \\ &\quad + \mathrm{dist}(w_2, v_j, G) \leq \mathrm{dist}(s_i, w_2, G) + 2 + \mathrm{dist}(w_2, v_j, G) \\ &= |\pi(s_i, v_j)| + 2 = \mathrm{dist}(s_i, v_j, G) + 2 \ .\end{aligned}$$

where the last inequality follows by the fact that $w \in C$. The theorem follows. ∎

Additive sourcewise spanners. The construction of additive sourcewise spanners combines the path-buying technique of [5,12,20] and the 4-additive spanner techniques of [11].

Theorem 13. *Let* $k \geq 1$ *be an integer. For every unweighted* n*-vertex graph* $G=(V,E)$ *and every subset* $S \subseteq V$*, there exists a (polynomially constructible)* $(2k, S)$*-additive sourcewise spanner* $H \subseteq G$ *of size* $\widetilde{O}(k \cdot n^{1+(k\varepsilon+1)/(2k+2)})$ *where* $\varepsilon = \log|S| / \log n$*.*

We first provide some high level overview of the algorithm. Inspired by [11], the algorithm handles separately "'distant" vertex pairs and "nearby" vertex pairs in $S \times V$, where the classification is based on some distance function $\mathrm{dist}^*(s, v, G)$.

To provide a bounded stretch for the "distant" vertex pairs (i.e., whose shortest path in G is long according to the function dist^*), the algorithm picks a small sample of vertices for rooting BFS trees. It is shown that this sample covers with high probability the neighborhood of the paths $\pi(s,v)$, when s and v are distant. Next, the path-buying procedure of [12] is applied on the collection of the "short" $s-v$ paths (i.e., handling the nearby vertex pairs in $S \times V$). In the analysis, we show that since the path-buying procedure cares only for the "short" $\pi(s,v)$ paths, the obtained spanner is sparser than that of [12].

The algorithm. We begin by briefly outline the strategy of the path buying procedure. In an initial clustering phase, a suitable clustering of the vertices is computed and an associated subset of edges is added to the spanner. This is followed by a path-buying phase, which examines certain paths in sequence, and decides whether or not to add them to the spanner. The decision is determined by assigning each candidate path a *cost*, corresponding to the number of path edges not already contained in the spanner, and a *value*, measuring how much adding the path would help to satisfy the considered set of constraints on the pairwise distances. The candidate path is added if the *value to cost ratio* is sufficiently large. The path-buying strategy was employed in the context of pairwise spanners both in [12] and [20].

The following notation is useful in our setting. A pair $(s_i, v_j) \in S \times V$ is *satisfied* by a subgraph $H \subseteq G$ if $\mathrm{dist}(s_i, v_j, H) \leq \mathrm{dist}(s_i, v_j, G) + 2k$. The neighborhood of a path P is denoted by $\Gamma(P) = \bigcup_{u \in P} \Gamma(u)$.

Define

$$Y = n^{(k\varepsilon+1)/(2k+2)} \quad \text{and} \quad L = n \log n / Y^2 \ . \tag{38}$$

A vertex v is *heavy* iff $\mathtt{deg}(v, G) \geq Y$, otherwise it is *light*. We now classify the $s_i - v_j$ paths $\pi(s_i, v_j)$ according to the number of heavy vertices that appear on the path. For a path P, let $\mathrm{dist}_{heavy}(P)$ be the number of heavy vertices in P. The path P is *long* if $\mathrm{dist}_{heavy}(P) \geq L$, otherwise it is *short*. The pair $\langle s_i, v_j \rangle \in S \times V$ is *long* (resp., *short*) iff $\pi(s_i, v_j)$ is long (resp., short). Define the subgraph

$$E(H_0) = \{(u,v) \in E(G) \ \mid \ \text{u is light } \} \ . \tag{39}$$

as the set of all edges adjacent to light vertices in G. The algorithm consists of two phases that add edges to H_0. The goal of the first phase is to satisfy the long pairs and the goal of the second phase is to satisfy the short pairs in the final spanner.

(1) Satisfying Long Pairs. Initially set $H^a = H_0$. Randomly select a set of vertices Z of expected size $9Y$, by choosing every vertex from V independently at random with probability $9Y/n$. For every vertex $z \in Z$, construct a BFS tree $BFS(z, G)$ rooted at z spanning all vertices V, and add the edges of $BFS(z, G)$ to H^a.

(2) Satisfying Short Pairs. This phase consists of two steps: a clustering step and a path-buying procedure.

(2.1) Clustering. Apply the clustering procedure of Fact 12 with $\gamma = \log_n Y$ which constructs a clustering $\mathcal{C} = \{C_1, \ldots, C_\lambda\}$ and a clustering subgraph G_C such that $|E(G_C)| = O(n \cdot Y)$ and $\lambda = O(n/Y)$.

(2.2) Path-buying. The following path-buying procedure is very similar to the construction of additive sourcewise spanner in [12]. The only modification is that in the current setting, the algorithm iterates over the *short* pairs $\langle s_i, v_j \rangle$ (as the long pairs are handled in the first phase) and not over all $S \times V$ pairs as in [12]. This fact enables the slight improvement in the size of the spanner.

Define $\mathcal{P}_{short} = \{\pi(s_i, v_j) \mid (s_i, v_j) \text{ is short }\}$ as the collection of short $s_i - v_j$ paths which are the candidates to be bought and added to the spanner and initialize $H_0^b \leftarrow G_C \cup H_0$. Iterate over all paths $\pi(s_i, v_j) \in \mathcal{P}_{short}$. At iteration $t \geq 0$, given the current spanner H_{t-1}^b, consider the t'th path in $\mathcal{P}_{short}$, let it be $\pi(s_i, v_j)$. Define paths $\pi_{i,j}^\ell$ for $0 \leq \ell \leq k$, maintaining the following invariants:
(i) $\pi_{i,j}^\ell$ is a path between s_i and v_j of length at most $\text{dist}(s_i, v_j, G) + 2\ell$.
(ii) any cluster $C \in \mathcal{C}$ contains at most three vertices of $\pi_{i,j}^\ell$,
(iii) $\texttt{Cost}(\pi_{i,j}^\ell) \leq L/\varphi^\ell$ where $\texttt{Cost}(\pi_{i,j}^\ell)$ is the number of edges of $\pi_{i,j}^\ell$ absent in the current spanner H_{t-1}^b and $\varphi = (2L)^{1/k}$.

In the analysis, we show that for every short pair (s_i, v_j), the algorithm buys exactly one path $\pi_{i,j}^\ell$ for $0 \leq \ell \leq k$, hence ensuring that the pair (s_i, v_j) is satisfied in H_t^b. We now describe the construction of the $\pi_{i,j}^\ell$ paths. For $\ell = 0$, define $\pi_{i,j}^0 = \pi(s_i, v_j)$, i.e., the shortest-path in G. Observe that for $j = 0$, Invariant (i) is trivially satisfied, Invariant (ii) is satisfied by Fact 12 (otherwise $\pi_{i,j}^0$ would not be a shortest-path), and Invariant (iii) is satisfied because the number of missing edges in H_0^b is at most L (i.e., the pair (s_i, v_j) is short). For a given path $\pi_{i,j}^\ell$, where $\ell \in \{0, \ldots, k\}$, define the function $\texttt{Val}(\pi_{i,j}^\ell)$ as the number of clusters $C \in \mathcal{C}$ such that there exists a vertex $v \in C \cap \pi_{i,j}^\ell$ and the distance between s_i and C in H_{t-1}^b is strictly greater then the distance between s_i and C in $\pi_{i,j}^\ell$, i.e., $\text{dist}(s_i, C, \pi_{i,j}^\ell) < \text{dist}(s_i, C, H_{t-1}^b)$. Formally,

$$\texttt{Val}(\pi_{i,j}^\ell) = |\{C \mid \exists v \in C \cap \pi_{i,j}^\ell \text{ and } \text{dist}(s_i, C, \pi_{i,j}^\ell) < \text{dist}(s_i, C, H_{t-1}^b)\}| \ .$$

The path $\pi_{i,j}^\ell$ is added to H_{t-1}^b resulting in H_t^b iff

$$\texttt{Cost}(\pi_{i,j}^\ell) \leq 3 \cdot \varphi \cdot \texttt{Val}(\pi_{i,j}^\ell).$$

In other words, if the condition above holds, then $H_t^b = H_{t-1}^b \cup \pi_{i,j}^\ell$, the remaining values of ℓ are ignored and the algorithm proceeds to the next short pair. Otherwise, the path $\pi_{i,j}^{\ell+1}$ is constructed as follows. Let R be the longest suffix of $\pi_{i,j}^\ell$ containing exactly $\lfloor \texttt{Cost}(\pi_{i,j}^\ell)/\varphi \rfloor$ edges the are missing in H_{t-1}^b. By the maximality of R, the edge e that precedes R is missing in H_{t-1}^b and hence by Fact 12 both endpoints of e (one of which is the first vertex of R) are clustered. Hence, at least $1 + \lfloor \texttt{Cost}(\pi_{i,j}^\ell)/\varphi \rfloor \geq \texttt{Cost}(\pi_{i,j}^\ell)/\varphi$ vertices of R are clustered. By Invariant (ii) there are at least $\texttt{Cost}(\pi_{i,j}^\ell)/(3\varphi)$ clusters in $\mathcal{C}$ having at least one vertex of R. Since $\pi_{i,j}^\ell$ was not bought, there exists a cluster $C \in \mathcal{C}$ containing

a vertex $x \in C$ of R such that the distance between s_i and C in H^b_{t-1} is at most the distance between s_i and x in $\pi^\ell_{i,j}$. The path $\pi^{\ell+1}_{i,j}$ is constructed by taking a shortest-path in H^b_{t-1} from s_i to the closest vertex $y \in C$, then add a path of length at most two between y and x (which exists since x and y belongs to the same cluster and by Fact 12), and finally add the suffix of R starting at x. We now claim that the path $\pi^\ell_{i,j}$ maintains the invariants. Invariant (i) follows immediately as the length of $\pi^\ell_{i,j}$ was extended by at most two edges (by the selection of y and the addition of the intercluster path of length two). In addition note that $\pi^{\ell+1}_{i,j}$ can be easily transformed into a path which is not longer and satisfies Invariant (ii). In particular as long as the current path contains at least four vertices of the same cluster $C' \in \mathcal{C}$, the path can be shorten as follows. Let a and b be the vertices of the current path closest to s_i and v_j respectively, replace the $a-b$ subpath (of length at least three) by an the intercluster $a-b$ path in G_C of length at most two. Consequently, Invariant (ii) is satisfied. Finally, note that the missing edges of $\pi^{\ell+1}_{i,j} \setminus H^b_{t-1}$ (i.e., the edges that contributes to $\texttt{Cost}(\pi^{\ell+1}_{i,j})$) are fully contained in R. By the definition of R, $|R| \leq \texttt{Cost}(\pi^\ell_{i,j})/\varphi^\ell \leq L/(\varphi^\ell \cdot \varphi)$, Invariant (iii) is satisfied. This completes the construction of $\pi^{\ell+1}_{i,j}$. Let $t' = |\mathcal{P}_{short}|$ be the number of paths considered to be bought. The $(2k, S)$-additive sourcewise spanner H is given by

$$H = H^a \cup H^b_{t'}.$$

Analysis. We begin by analyzing the first phase of the algorithm and show that the long pairs are satisfied in H^a.

Lemma 10. *(1) With probability at least $1-1/n$, every long pair $(s_i, v_j) \in S \times V$ is satisfied in H^a. (2) H^a has $\widetilde{O}(n^{1+(k\varepsilon+1)/(2k+2)})$ edges in expectation.*

Proof: Begin with part (1) and consider some long $s_i - v_j$ path $\pi(s_i, v_j)$. Since $\pi(s_i, v_j)$ is a shortest-path in G, every vertex $v \in \pi(s_i, v_j)$ has at most two neighbors in $\pi(s_i, v_j)$. Combining this with the fact that $\pi(s_i, v_j)$ contains at least L heavy vertices whose sum of degrees is more than $Y \cdot L$, we get that $|\Gamma(\pi(s_i, v_j))| \geq Y \cdot L/3$. We now claim that the probability that $\Gamma(\pi(s_i, v_j)) \cap Z \neq \emptyset$ is at least $1 - 1/n^3$, as

$$\mathbb{P}[\Gamma(\pi(s_i, v_j)) \cap Z = \emptyset] \leq (1 - 9Y/n)^{Y \cdot L/3} \leq 1/n^3,$$

which follows by Eq. (38). Part (1) of the lemma follows by applying the union bound over all pairs of vertices (although it is sufficient to consider only the $S \times V$ pairs). Consider part (2). The size of H_0 is $O(nY)$ as it consists of edges adjacent to light vertices. The expected number of vertices in Z is $O(Y)$. For each $z \in Z$, a BFS tree of $n-1$ edges is added to H. Hence, the expected number of edges in H^a is $O(n \cdot Y)$. Part (2) follows by plugging Eq. (38). ∎

We now turn to consider the short pairs and analyze the second phase of the algorithm.

Lemma 11. *(1) Every short pair $(s_i, v_j) \in S \times V$ is satisfied in H^b. (2) H^b contains $\widetilde{O}(k \cdot n^{1+(k\varepsilon+1)/(2k+2)})$ edges.*

Proof: Begin with (1). By Invariant (iii) of the path-buying procedure, we have that $\texttt{Cost}(\pi^k_{i,j}) \leq 1/2$, since the $\texttt{Cost}()$ function has only integral values, it has to be that $\texttt{Cost}(\pi^k_{i,j}) = 0$, which ensures that $\pi^\ell_{i,j}$ is bought for some $\ell \leq k$ for every $\pi^0_{i,j} \in \mathcal{P}_{short}$. By the above, for every short pair (s_i, v_j), there exists some $\ell \in \{0, \ldots, k\}$ such that $\pi^\ell_{i,j}$ was bought, hence the claim holds by Invariant (i). Consider part (2). By Fact 12, G_C contains $n \cdot Y$ edges. Let $\mathcal{B}$ be the paths bought during the path-buying procedure. It remains to bound the number of edges added due to the paths in $\mathcal{B}$. For every short pair (s_i, v_j) let $P_{i,j}$ be the path in $\mathcal{B}$, i.e., $P_{i,j} = \pi^\ell_{i,j}$ for some $\ell \in \{0, \ldots, k\}$ (by Part (1) above such ℓ exists). We first claim that very cluster $C \in \mathcal{C}$ contributes to $\texttt{Val}(P_{i,j})$ of at most $|S|(2k+3)$ bought paths. This holds since when for $s_i \in S$ a supported path is bought the distance between s_i and C is at most $2k+2$ greater than the distance between s_i and C in G: otherwise one could shorten $P_{i,j}$ by more than $2k$, obtaining a contradiction with Invariant (i). Therefore the total number of edges added during the path buying procedure is upper bounded by $\sum_{i,j} \texttt{Cost}(P_{i,j}) \leq \sum_{i,j} 3\varphi \cdot \texttt{Val}(P_{i,j}) \leq 3\varphi \cdot (2k+3) \cdot |S| \cdot n/Y$. By plugging Eq. (38) and recalling that $|S| = O(n^\varepsilon)$, the lemma follows. ∎

Theorem 13 follows by Lemma 10 and 11.

Finally, we provide an "almost" tight construction for $(4, S)$-sourcewise additive spanners for a sufficiently large subset of sources S. We use the following fact due to [12].

Fact 14 (Additive subsetwise spanners [12]) *For every unweighted n-vertex graph $G = (V, E)$, and a subset $Z \subseteq V$, there exists a (polynomially constructible) subgraph $H \subseteq G$ with $O(n^{1+\kappa/2})$ edges such that $\text{dist}(u, v, H) \leq \text{dist}(u, v, G) + 2$ for every $(u, v) \in Z \times Z$, where $\kappa = \log |Z| / \log n$ (this subgraph is also known as* additive subsetwise spanner*).*

We have the following.

Theorem 15. *For every unweighted n-vertex graph $G = (V, E)$ and a subset of sources $S \subseteq V$ such that $|S| = \Omega(n^{2/3})$, there exists a (polynomially constructible) $(4, S)$-additive sourcewise spanner $H \subseteq G$ with $O(n^{1+\varepsilon/2})$ edges.*

Proof: The algorithm is as follows. (i) Apply the clustering procedure of Fact 12 with $\gamma = \varepsilon/2$. This constructs a clustering $\mathcal{C} = \{C_1, \ldots, C_\lambda\}$ and a clustering subgraph G_C such that $|E(G_C)| = O(n^{1+\varepsilon/2})$ and $\lambda = O(n^{1-\varepsilon/2})$. Let $X = \{x_1, \ldots, x_\lambda\}$ be the cluster centers of $\mathcal{C}$. (ii) Set $Z = X \cup S$ and apply the 2-additive subsetwise spanner construction of Fact 14 on Z, resulting in the subgraph H'. Let $H = G_C \cup H'$. Note that $|Z| = O(n^\varepsilon)$, since $\varepsilon \geq 2/3$. Hence by Facts 12 and 14, it holds that H has $O(n^{1+\varepsilon/2})$ edges. We now analyze the $S \times V$ stretch in H. Consider some $\pi(s_i, v_j)$ path that has some missing edge in G_C, i.e., $\pi(s_i, v_j) \setminus G_C \neq \emptyset$. Let $e = (y_1, y_2)$ be the last missing edge on the path (closest to v_j). By Fact 12, y_2 is clustered. Let x be the cluster center of $C(y_2)$.

Hence, $x \in Z$. Define $P_1 \in SP(s_i, x, H)$, $P_2 = (x, y_2)$ and $P_3 = \pi(s_i, v_j)[y_2, v_j]$. Let $P = P_1 \circ P_2 \circ P_3$. It then holds that $P \subseteq H$. By step (ii) it holds that $|P_1| \leq \text{dist}(s_i, x, G) + 2$. We have the following.

$$\begin{aligned}\text{dist}(s_i, v_j, H) &\leq |P| = |P_1| + |P_2| + |P_3| \\ &\leq \text{dist}(s_i, x, G) + 2 + 1 + \text{dist}(y_2, v_j, G) \\ &\leq \text{dist}(s_i, y_2, G) + 1 + 3 + \text{dist}(y_2, v_j, G) = |\pi(s_i, v_j)| + 4 \\ &= \text{dist}(s_i, v_j, G) + 4 \ ,\end{aligned}$$

where the last inequality follows by the fact the w is the cluster center of y_2 and by the triangle inequality. ∎

4 Conclusion

In this paper, we considered the following question: what is the "minimal" modification to the standard definition of $(2k-1)$-spanners required in order to bypass the barrier of $\Omega(n^{1+1/k})$ edges, for every integer $k \geq 1$. We proposed two such modifications. First, we make a hard distinction between pairs at distance 1 and pairs at distance larger than 1, corresponding to the hard distinction arising from lower bounds that are based on the girth conjecture. This results in a k-hybrid spanner with $O(n^{1+1/(k+1)})$ edges. The k-hybrid spanner can be viewed either as a *relaxation* for k-spanners (i.e., relaxing the requirement of multiplicative stretch k for neighboring vertex pairs), or alternatively as a *strengthening* of (α, β)-spanners. Whereas (α, β)-spanners make a somewhat "fuzzy" distinction between large and small distances, hybrid spanners provide a *sharp* distinction between vertex pairs at distance 1 and pairs at distance larger than 1, which corresponds to the sharp distinction arising by the girth argument. In the second modification, we considered sourcewise spanners, where we care only for the stretch of vertex pairs from $S \times V$ for a given subset of sources $S \subseteq V$. In particular, it follows from our results that even if one considers a large subset of sources $S \subseteq V$ for $|S| = O(n^{1-\delta})$ for sufficiently small, but fixed δ, there exist a subgraph with $O(n^{1+(1-\delta)/k})$ edges providing a multiplicative stretch of $2k-1$ for every pair $(u, v) \in S \times V$. An interesting future direction is to study the lower bounds for hybrid spanners and sourcewise multiplicative spanners.

Acknowledgment. I am very grateful to my advisor, Prof. David Peleg, for many helpful discussions and for reviewing this paper. I would also like to thank Michael Dinitz and Eylon Yogev for useful comments and discussions.

References

1. R. Agarwal and P.B. Godfrey and S. Har-Peled. Approximate distance queries and compact routing in sparse graphs. In *Proc. INFOCOM*, 2011.

2. D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.
3. N. Alon and J.H. Spencer. The probabilistic method. *Wiley*, New York, 1992.
4. I. Althöfer, G. Das, D. Dobkin, D. Joseph and J. Soares. On sparse spanners of weighted graphs. *Networks*, 9(1): 81–100, 1993.
5. S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie. Additive spanners and (α, β)-spanners. *ACM Trans. Algo.* 7, A.5, 2010.
6. S. Baswana and S. Sen. A simple Linear Time Randomized Algorithm for Computing Sparse Spanners in Weighted Graphs. In *Random Structures and Algorithms*, 30(4): 532–563, 2007.
7. S. Baswana and T. Kavitha. Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In *Proc. FOCS*, 591–602, 2006.
8. S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in expected $O(n^2)$ time. *ACM Transactions on Algorithms (TALG)*, 2(4):557–577, 2006.
9. B. Bollobás, D. Coppersmith and M. Elkin. Sparse distance preservers and additive spanners. *SIAM Journal on Discrete Mathematics*, 19(4): 1029–1055, 2005.
10. D. Coppersmith, M. Elkin. Sparse sourcewise and pairwise distance preservers. In *SIAM Journal on Discrete Mathematics*, 20(2): 463–501, 2006.
11. S. Chechik. New Additive Spanners. In *Proc. SODA*, 29(5): 498–512, 2013.
12. M. Cygan, F. Grandoni and T. Kavitha. On Pairwise Spanners. In *Proc. STACS*, 209–220, 2013.
13. C. Gavoille and D. Peleg. Compact and localized distributed data structures. *Distributed Computing*, 16(2): 111-120, 2003.
14. D. Dor, S. Halperin and U. Zwick. All-pairs almost shortest paths. *SIAM on Computing*, 29(5): 1740–1759, 2000.
15. M. Elkin and D. Peleg. $(1+\varepsilon, \beta)$-Spanner Constructions for General Graphs. *SIAM Journal on Computing*, 33(3): 608–631, 2004.
16. M. Elkin. Computing almost shortest paths. *ACM Transactions on Algorithms (TALG)*, 1(2): 283–323, 2005.
17. M. Elkin, Y. Emek, D.A. Spielman and S.H. Teng. Lower stretch spanning trees. In *Proc. STOC*, 494–503, 2005.
18. P. Erdős. *Extremal problems in graph theory.* In *Proc. Symp. Theory of Graphs and its Applications*, page 2936, 1963.
19. M. Kapralov and R. Panigrahy. Spectral sparsification via random spanners. *ITCS*, 2012.
20. T. Kavitha and N.M Varma. Small Stretch Pairwise Spanners. In *ICALP*, 601–612, 2013.
21. A.L Liestman and T.C. Shermer. Additive graph spanners. *Networks*, 23(4): 343–363, 1993.
22. M. Mendel and A. Naor. Ramsey partitions and proximity data structures. *FOCS*, 23(4): 109–118, 2006.
23. M. Pătraşcu and L. Roditty. Distance oracles beyond the Thorup-Zwick bound. In *FOCS*, 815–823, 2010.
24. D. Peleg and A.A Schaffer. Graph spanners. In *Journal of graph theory*, 12(1):99-116, 1989.
25. D. Peleg, J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM Journal on computing*, 18(4): 740–747, 1989.
26. D. Peleg. *Distributed Computing: A Locality-Sensitive Approach.* SIAM, 2000.

27. S. Pettie. Low distortion spanners. In *ACM Transactions on Algorithms (TALG)*, 6(1), 2009.
28. L. Roditty, M. Thorup and U. Zwick. Deterministic constructions of approximate distance oracles and spanners. In *ICALP*, 261–272, 2005.
29. M. Thorup. Undirected single-source shortest paths with positive integer weights in linear time. In *Journal of the ACM (JACM)*, 46(3): 362–394, 1999.
30. M. Thorup and U. Zwick. Approximate distance oracles. In *Journal of the ACM (JACM)*, 52(1):1–24, 2005.
31. M. Thorup, U. Zwick. Spanners and emulators with sublinear distance errors. In *SODA*, 802–809, 2006.
32. R. Wenger. Extremal graphs with no C4's, C6's, or C10's. *Journal of Combinatorial Theory*, 113–116, 1991.
33. D.P Woodruff. Lower bounds for additive spanners, emulators, and more. In *Proc. 47th Symp. on Foundations of Computer Science*, 389–398, 2006.